arXiv:1510.03964v1 [q-bio.GN] 14 Oct 2015

Pathway Tools version 19.0: Integrated Software for Pathway/Genome Informatics and Systems Biology

Peter D. Karp*, Mario Latendresse, Suzanne M. Paley, Markus Krummenacker Quang Ong, Richard Billington, Anamika Kothari, Daniel Weaver, Tom Lee, Pallavi Subhraveti, Aaron Spaulding, Carol Fulcher, Ingrid M. Keseler, and Ron Caspi

Bioinformatics Research Group, SRI International
333 Ravenswood Ave, Menlo Park, CA 94025
pkarp@ai.sri.com

*To whom correspondence should be addressed.

Abstract

Pathway Tools is a bioinformatics software environment with a broad set of capabilities. The software provides genome-informatics tools such as a genome browser, sequence alignments, a genome-variant analyzer, and comparative-genomics operations. It offers metabolic-informatics tools, such as metabolic reconstruction, quantitative metabolic modeling, prediction of reaction atom mappings, and metabolic route search. Pathway Tools also provides regulatory-informatics tools, such as the ability to represent and visualize a wide range of regulatory interactions. The software creates and manages a type of organism-specific database called a Pathway/Genome Database (PGDB), which the software enables database curators to interactively edit. It supports web publishing of PGDBs and provides a large number of query, visualization, and omics-data analysis tools. Scientists around the world have created more than 9,800 PGDBs by using Pathway Tools, many of which are curated databases for important model organisms. Those PGDBs can be exchanged using a peer-to-peer database-sharing system called the PGDB Registry.

Biographical Note

Dr. Peter Karp is the Director of the Bioinformatics Research Group at SRI International. He received the PhD degree in Computer Science from Stanford University.

1 Introduction

Pathway Tools [1, 2, 3, 4] is a software environment for management, analysis, simulation, and visualization of integrated collections of genome, pathway, and regulatory data. This article provides an overview of all current capabilities of Pathway Tools. A shorter version of this article describing only those capabilities of Pathway Tools developed since publication of a 2010 article on the software [1] was recently published in *Briefings in Bioinformatics* [5]. New software capabilities developed since 2010 are typeset in blue in this article.[1]

Pathway Tools handles many types of information beyond pathways and offers extensive capabilities. The software has been under continuous development within the Bioinformatics Research Group (BRG) within SRI International since the early 1990s. Pathway Tools serves several different use cases in bioinformatics and systems biology:

- It supports development of organism-specific databases (also called model-organism databases) that integrate many bioinformatics data types.
- It supports scientific visualization, web publishing, and dissemination of those organism-specific databases.
- It performs computational inferences from sequenced genomes including prediction of an organism's metabolic network, prediction of metabolic pathway hole fillers, and prediction of operons.
- It enables creation of steady-state quantitative metabolic flux models for individual organisms and for organism communities.
- It provides tools for graph-based analysis of biological networks, such as for identification of metabolic choke points, dead-end metabolites, and blocked reactions.
- It provides tools for analysis of gene expression, metabolomics, proteomics, and multi-omics datasets.
- It provides comparative analyses of organism-specific databases.
- It supports metabolic engineering.

Pathway Tools is focused around a type of model-organism database called a Pathway/Genome Database (PGDB). A PGDB integrates information about an organism's genes, proteins, metabolic network, and regulatory network.

Pathway Tools has several components. The **PathoLogic** component enables users to create a new PGDB from the annotated genome of an organism, containing the genes, proteins, biochemical reactions, and predicted metabolic pathways and operons of the organism.

[1] This article incorporates significant text from [1] by permission of the publisher.

The **Pathway/Genome Editors** let PGDB developers interactively refine the contents of a PGDB, such as editing a metabolic pathway or an operon, or defining the function of a newly characterized gene.

The **Pathway/Genome Navigator** supports querying, visualization, and analysis of PGDBs. Whereas all other Pathway Tools components run as desktop applications only, the Navigator can run as both a desktop application and as a web server. The Navigator enables scientists to quickly find information, to display that information in familiar graphical forms, and to publish a PGDB to the scientific community via the web. The Navigator provides a platform for systems-level analysis of high-throughput data by providing tools for painting combinations of gene expression, protein expression, and metabolomics data onto a full metabolic map of the cell, onto the full genome, and onto a diagram of the regulatory network of the cell.

The **MetaFlux** component enables construction and execution of steady-state metabolic flux models from PGDBs. MetaFlux has modes to accelerate development of metabolic models, and to use metabolic models to simulate both gene and reaction knockouts. Pathway Tools provides a unique environment for metabolic flux modeling: by combining a tool for reconstructing metabolic networks from genome annotations with metabolic-model debugging tools such as a reaction gap-filler, the software enables rapid development of metabolic models from sequenced genomes. And by tightly coupling the metabolic model with other enriching information such as the sequenced genome, chemical structures, and regulatory information, Pathway Tools-based metabolic models are easier to understand, validate, reuse, extend, and learn from.

Pathway Tools includes a sophisticated ontology and database application programming interface (API) that enables programs to perform complex queries, symbolic computations, and data mining on the contents of a PGDB. For example, the software has been used for global studies of the *E. coli* metabolic network [6] and genetic network [7].

Pathway Tools is seeing widespread use across the bioinformatics community to create Pathway/Genome Databases in all domains of life. More than 5,600 groups to date have licensed the software. As well as supporting the development of the EcoCyc [8] and MetaCyc [9] databases (DBs) at SRI, and SRI's BioCyc collection of 5,500 PGDBs [9], the software is used by genome centers, experimental biologists, and groups that are creating curated DBs for a number of different organisms (see Section 11 for a more detailed listing of available PGDBs).

This article provides a comprehensive description of Pathway Tools. Where possible, it references earlier publications that provide more algorithmic details. However, in some cases, those earlier publications are outdated by new developments in the software that are described here. This article also emphasizes new aspects of the software that have not been reported in earlier publications. The organization of this article is as follows. Section 2 articulates in more detail the use cases for which Pathway Tools was designed. Section 3

relates how a new PGDB is created, and describes the computational inference procedures within Pathway Tools. It summarizes Pathway Tools' interactive editing capabilities and the associated author-crediting system. It also describes tools for automatic upgrading of a PGDB schema, and for bulk updating of the genome annotation within a PGDB. Section 4 describes the schema of a PGDB and the ontologies used by Pathway Tools. Section 5 relates the querying and visualization facilities of Pathway Tools. Section 6 summarizes the mechanisms for importing and exporting data from Pathway Tools, and for accessing and updating PGDB data via APIs. Section 7 describes multiple Pathway Tools modules for performing network analyses of PGDBs, such as a tool for identifying dead-end metabolites. Section 8 describes comparative analysis tools within the software. Section 9 describes the software architecture of Pathway Tools. Section 11 lists the large family of PGDBs that have been created by Pathway Tools users outside SRI International, and describes a peer-to-peer data-sharing facility within Pathway Tools that enables users to easily exchange their PGDBs. Section 10 describes the metabolic modeling capabilities of the software. Section 12 compares Pathway Tools to related efforts.

2 Pathway Tools Use Cases

This section articulates the objectives for which Pathway Tools was designed. Please note that when we assert that Pathway Tools supports a given type of use case, it does not mean that Pathway Tools provides every type of computational tool needed in that area. For example, omics data analysis is a huge field, and although Pathway Tools contributes novel and useful omics data analysis capabilities, it does not provide every omics data analysis method: in fact, it is intended to be used in conjunction with other omics analysis tools (such as for data normalization). Section 13 summarizes the limitations of Pathway Tools.

2.1 Development of Organism-Specific Databases

Organism-specific DBs (also known as model-organism DBs) describe the genome and other information about an organism [8, 10, 11, 12, 13]. We posit that every organism with a completely sequenced genome and an experimental community of significant size requires an organism-specific DB to fully exploit the genome sequence. Such DBs should provide a central information resource that integrates information dispersed through the scientific literature about the genome, molecular parts, and cellular networks of the organism. Such DBs both direct and accelerate further scientific investigations.

Pathway Tools facilitates rapid initial computational construction of organism-specific DBs, followed by manual refinement of the PGDB, to produce an extremely rich and accurate DB in minimal time. Our approach differentiates experimental versus computationally inferred information whenever possible. Rapid construction of PGDBs is achieved by importing an

annotated genome into a PGDB in the form of a GenBank file, and by applying several computational inference tools to infer new information within the PGDB, such as metabolic pathways. Scientists can then employ the Pathway/Genome Editors to correct and supplement computational inferences when necessary, and to perform ongoing manual curation of the PGDB if desired.

The Pathway Tools DB schema (see Section 4) is significant in both its breadth and its depth: it models an unusually broad set of bioinformatics data types ranging from genomes to pathways to regulatory networks, and it provides high-fidelity representations of those data types that allow PGDBs to accurately capture complex biology.

2.2 Web Visualization and Querying of Organism-Specific Databases

To speed user comprehension of the complex information within PGDBs, the Pathway/Genome Navigator provides many scientific visualization services, including a genome browser, visualization of single metabolic pathways and entire metabolic maps, visualization of single operons and of entire regulatory networks, and visualization of chemical compounds and reactions (see Section 5 for more details). These visualization tools operate within a web server, permitting developers of PGDBs to publish their PGDBs to the scientific community through a website. This form of PGDB publishing supports interactive querying and browsing using a three-tiered series of web query interfaces (see Section 5.1), including a quick search, a set of object-specific query tools, and a tool for interactively constructing queries whose power is comparable to that of SQL.

We have developed other publishing paradigms to support computational analysis and dissemination of PGDBs. Pathway Tools APIs exist in four languages [14] and as web services. PGDBs can be exported in several formats and imported into the BioWarehouse DB integration system [15]. Finally, users can easily share and exchange PGDBs using a peer-to-peer DB-sharing system that we have developed.

2.3 Extend Genome Annotations with Additional Computational Inferences

Pathway Tools extends the paradigm of genome analysis. After traditional analyses such as gene calling and gene function prediction are performed by external software packages, Pathway Tools provides additional computational genome analyses that layer additional information above the traditional genome annotation. Pathway Tools predicts the operons of the organism. It predicts the metabolic pathways of the organism. It also predicts which genes in the organism code for missing enzymes in the predicted metabolic pathways, thus

using pathway information to predict additional gene functions. See Section 3 for more details.

2.4 Analysis of Omics Data

Pathway Tools provides three genome-scale viewers for animated visualization of omics datasets in the context of the full metabolic network [2], full transcriptional regulatory network, and full genome (see Sections 5.4–5.6 for more details). It also provides enrichment analysis and SmartTable-based analysis of omics data.

2.5 Quantitative Metabolic Flux Modeling

The MetaFlux module of Pathway Tools supports development and execution of steady-state metabolic flux models for individual organisms and organism communities from PGDBs. MetaFlux supports a *literate modeling* approach that makes metabolic flux models highly accessible to and understandable by scientists.

2.6 Analysis of Biological Networks

Pathway Tools includes programs for symbolic analysis of biological networks (see Section 7 for more details) that rely on the detailed biological network ontology underlying Pathway Tools. The software identifies dead-end metabolites and blocked reactions, both of which usually reflect errors or incompleteness of our knowledge of a metabolic network.

Pathway Tools indirectly supports a two-phased, pathway-based paradigm for drug discovery. Phase I is the search for essential in vivo metabolic pathways: pathways whose function is essential for microbial growth in the host. Phase II is the search for targets within essential in vivo pathways. Both phases are supported by a Pathway Tools module that predicts choke-point reactions within the metabolic network as likely drug targets [16].

2.7 Comparative Analyses of Organism-Specific Databases

Pathway Tools provides a suite of comparative analysis operations that can be applied to multiple user-selected PGDBs (see Section 8 for more details). Pathway Tools emphasizes comparisons at the functional level, rather than the sequence level. Example comparisons include (1) highlighting on the Cellular Overview of one organism the reactions that it shares (or does not share) with one or more other organisms; (2) a tabular comparison of the reaction complements of several organisms, organized by substrate type (e.g., small molecules, RNAs, proteins) or by the number of isozymes per reaction; (3) a comparison of

the pathway complements of several organisms, where the tabular pathway comparison is organized by a pathway ontology; (4) a table showing which genes have orthologs in which PGDBs; and (5) a comparison of the genome organization of orthologs using the genome browser.

2.8 Metabolic Engineering

Metabolic engineering is a discipline that seeks to modify the metabolic network of an organism in a desired fashion, such as to achieve overproduction of desired end products, or degradation of specified compounds [17]. Pathway Tools is designed to assist metabolic engineers in several respects. Its metabolic-reconstruction capabilities aid in rapid characterization of a host organism for metabolic engineering. Its editing tools permit refinement of that pathway database. Its omics analysis capabilities aid metabolic engineers in understanding the activity levels of different portions of the metabolic network under different growth conditions. Its RouteSearch tool supports design of novel reaction pathways from a feedstock compound to a desired product compound, and its metabolic modeling capabilities enable computational exploration of modified flux routes.

3 Creating and Curating a PGDB

The life cycle of a PGDB typically includes the following three procedures.

1. Initial creation of the PGDB starts with one or more input files describing the functionally annotated genome of an organism. The PathoLogic component of Pathway Tools transforms the genome into an Ocelot [18] DB structured according to the Pathway Tools schema. Next, the user applies one or more computational inference tools within PathoLogic to the genome to infer new information such as metabolic pathways. For several of the PathoLogic inference tools, we have created graphical user interfaces (GUIs) that enable the user to review the inferences made by these tools, and to accept, reject, or modify those inferences.

2. PGDB curation. Manual refinement and updating of a PGDB is performed using the Pathway/Genome Editors. This phase can last for years, or for decades, as in the case of EcoCyc [8]. Curation can be based on information found about the organism in the experimental literature, information from in-house experiments, or information inferred by the curator, perhaps with help from computational tools. PGDB curation is multidimensional, involving addition and/or deletion of genes or metabolic pathways to/from the PGDB; changing gene functions; altering the structure of metabolic pathways; authoring of summary comments for genes or pathways; attachment of Gene Ontology terms to genes and gene products; entry of chemical structures for small molecules; defining regulatory relationships; and entry of

data into many different PGDB fields including protein molecular weights, pIs, and cellular locations.

3. Bulk updating of a PGDB. A PGDB developer might run an external program that predicts cellular locations for hundreds of genes, and want to load those predictions into the PGDB. Some groups store the authoritative genome annotation for an organism in another genome-data management system, and want to periodically import the latest genome annotation into Pathway Tools; the software provides a tool for this operation. In addition, some of the individual components within PathoLogic that were used to initially create a PGDB can be run again at a later date to take advantage of updated information.

The following subsections describe these procedures in more detail.

3.1 PathoLogic PGDB Creation

PathoLogic performs a series of computational inferences that are summarized in Figure 1. These inferences can be performed in an interactive mode, in which the user guides the system through each step, and can review and modify the inferences made by the system using interactive tools. PathoLogic can also execute in a batch mode, in which all processing is automated, to process hundreds or thousands of genomes.

The input to PathoLogic is the annotated genome of an organism. PathoLogic does not perform genome annotation; its input must supply the genome sequence, the locations of genes, and functions of gene products. The annotation is supplied as a set of files in GenBank format or PathoLogic format, each of which describes the annotation of one replicon (chromosome or plasmid), or of one contig for genomes that are not fully assembled. When the annotation is provided in PathoLogic format, the sequence is provided as one or more separate FASTA files.

The annotation specified in a GenBank or PathoLogic file can include the start and stop positions of the coding region for each gene, and intron positions. It can also include a description of the function of the gene product as a text string, one or more EC numbers, and one or more Gene Ontology terms. The annotation can also include a gene name, synonyms for the gene name and the product name, links to other bioinformatics databases, and comments.

PathoLogic initializes the schema of the new PGDB by copying from MetaCyc into the new PGDB the definitions of the ontology classes and the 350 slots (DB attributes) that define the schema of a PGDB.

PathoLogic next creates a PGDB object for every replicon and contig defined by the input files, and for every gene and gene product defined in the input files. It populates these new objects with data from the input files, such as gene names and their sequence coordinates, and gene product names. As a result of these operations, the new PGDB now mirrors the

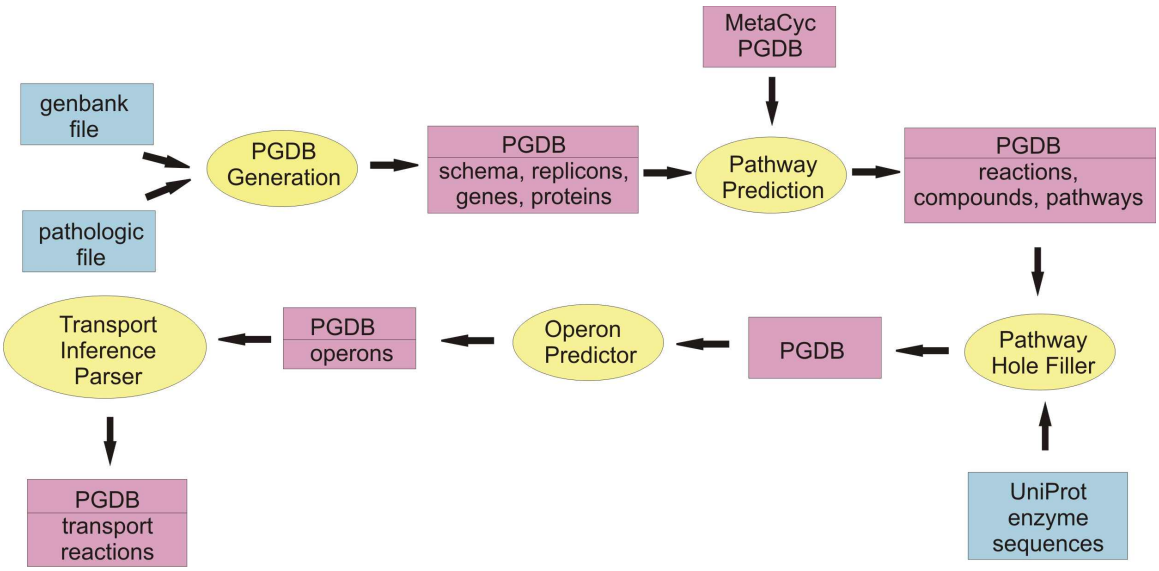

Figure 1: Inputs and outputs of the computational inference modules within PathoLogic. The initial input to PathoLogic is either a GenBank-format file or a PathoLogic-format file. The pink boxes all indicate that a PGDB is an input to or an output from some processing step; the notations at the bottom of the pink boxes indicate what types of data have been added by the previous processing step (for example, the Transport Inference Parser adds transport reactions to a PGDB).

information in the input files.

3.2 PathoLogic Inference of Metabolic Pathways

Pathway Tools predicts the metabolic pathway complement of an organism by assessing what known pathways from the MetaCyc PGDB [9] are present in the annotated genome of that organism's PGDB. This inference is performed in two steps that are described and evaluated in [19, 20, 21].

Step 1: The enzymes in the PGDB are assigned to their corresponding reactions in MetaCyc, thus defining the reactome of the organism. PathoLogic performs this assignment by matching to MetaCyc reactions the gene-product names (enzyme names), the EC numbers, and the Gene Ontology terms assigned to genes in the genome. The program can use whatever combination of these three information types is available in a given genome. For example, the *fabD* gene in *Bacillus anthracis* was annotated with the function "malonyl CoA-acyl carrier protein transacylase." That name was recognized by PathoLogic as corresponding to the MetaCyc reaction whose EC number is 2.3.1.39. PathoLogic therefore imported that reaction and its substrates into the *B. anthracis* PGDB, and created an enzymatic-reaction object linking that reaction to that *B. anthracis* protein.

Step 2: Once the reactome of the organism has been established in the preceding manner, PathoLogic imports into the new PGDB all MetaCyc pathways that contain at least one reaction in the organism's reactome. Once imported, PathoLogic then attempts to prune out those pathways that are likely to be false-positive predictions. That pruning process considers both the fraction of reaction steps in the pathway that have assigned enzymes and how many of the reactions with assigned enzymes are unique to that pathway (as opposed to being used in additional metabolic pathways in that organism). The remaining pathways are those predicted to occur in the organism under analysis. MetaCyc records curated information about the expected taxonomic groups in which a pathway is expected to occur (e.g., in plants only) based on experimental observations of that pathway to date. The software prunes a predicted pathway from organism X if organism X is outside the expected taxonomic distribution of that pathway.

3.2.1 Calculation of Pathway Abundance for Metagenomics Analysis

PathoLogic computes abundances of metabolic pathways based on gene abundances, which is useful for comparing the metabolic profiles of different microbial communities. Gene abundances are specified in the annotated genome file (PathoLogic format only).

No pre-processing of the gene abundances (such as outlier removal) is done by PathoLogic. The abundance of a pathway is computed based on the gene abundances involved in the pathway. More precisely, assume that R is the set of reactions in pathway P for which gene abundances are specified, $|R|$ is the size of R, and g_a is the given abundance of gene g. The abundance of a pathway P is

$$\sum_{r \in P} r_a / |R| \text{ where } r_a = \sum_{g \text{catalyzes } r} g_a$$

That is, the abundance of a pathway is the sum of the abundances of the genes catalyzing the reactions of the pathway, divided by the number of reactions of the pathway for which gene abundances are given. Notice that this formula does take into account all the known isozymes catalyzing a reaction and the spontaneous reactions do not take part in the computation. The abundances are provided, among other results, in the file `pathways-report.txt`.

3.3 PathoLogic Inference of Operons

The Pathway Tools operon predictor identifies operon boundaries by examining pairs of adjacent genes A and B and using information such as intergenic distance, and whether it can identify a functional relationship between A and B, such as membership in the same pathway [22], membership in the same multimeric protein complex, or whether A is a transporter for a substrate within a metabolic pathway in which B is an enzyme.

3.4 PathoLogic Inference of Pathway Holes

A pathway hole is a reaction in a metabolic pathway for which no enzyme has been identified in the genome that catalyzes that reaction. Typical microbial genomes contain 200–300 pathway holes. Although some pathway holes are probably genuine, we believe that the majority are likely to result from the failure of the genome-annotation process to identify the genes corresponding to those pathway holes. For example, genome-annotation systems systematically under-annotate genes with multiple functions, and we believe that the enzyme functions for many pathway holes are unidentified second functions for genes that have one assigned function. Erroneous pathway holes can result from the prediction of pathways that are not actually present in the organism.

The pathway hole-filling program PHFiller [23] (a component of PathoLogic) generates hypotheses as to which genes code for these missing enzymes by using the following method. Given a reaction that is a pathway hole, the program first queries the UniProt database to find all known sequences for enzymes that catalyze that same reaction in other organisms. The program then uses the BLAST tool to compare that set of sequences against the full proteome of the organism in which we are seeking hole fillers. It scores the resulting BLAST hits using a Bayesian classifier that considers information such as genome localization (that is, is a potential hole filler in the same operon as another gene in the same metabolic pathway?). At a stringent probability-score cutoff, our method finds potential hole fillers for approximately 45% of the pathway holes in a microbial genome [23].

PHFiller includes a graphical interface that optionally presents each inferred hole filler to the user along with information that helps the user evaluate the hole fillers, and enables the user to accept or reject the hole fillers that it has proposed.

3.5 PathoLogic Inference of Transport Reactions

Membrane transport proteins typically make up 5–15% of the gene content of organisms sequenced to date. Transporters import nutrients into the cell, thus determining the environments in which cell growth is possible. The development of the PathoLogic Transport Inference Parser (TIP) [24] was motivated by the need to perform symbolic inferences on cellular transport systems, and by the need to include transporters on the Cellular Overview diagram. The motivating symbolic inferences include the problems of computing answers to the following queries: What chemicals can the organism import or export? For which cellular metabolites that are consumed by metabolic reactions but never produced by any reaction does no known transporter exist (meaning that the origin of such metabolites is a mystery, and indicates missing knowledge about transporters or reactions that produce the compound)?

To answer such queries, Pathway Tools uses an ontology-based representation of transporter function in which transport events are represented as reactions in which the transported

compound(s) are substrates. Each substrate is labeled with the cellular compartment in which it resides, and each substrate is a controlled-vocabulary term from the extensive set of chemical compounds in MetaCyc [9]. The TIP program converts the free-text descriptions of transporter functions found in genome annotations (examples: "phosphate ABC transporter" and "sodium/proline symporter") into computable transport reactions.

3.6 Atom Mappings

The atom mapping of a reaction specifies for each reactant non-hydrogen atom its corresponding atom in a product compound. Pathway Tools contains an algorithm for computing atom mappings, described in [25]. Essentially, this approach computes atom mappings that minimize the overall cost of bonds broken and made in the reaction, given assigned propensities for bond creation and breakage. This algorithm has been applied to compute atom mappings for almost all of the reactions in the MetaCyc database.

Atom mappings are used in two other parts of Pathway Tools. Atom mappings are used in the rendering of Pathway Tools reaction pages, to depict the conserved chemical moieties in a reaction. Conserved moieties are depicted by using the same color on the reactant and product sides. The bonds made or broken by a reaction are identified from the atom mapping for the reaction, and are colored black. Atom mappings are also used in the RouteSearch module of Pathway Tools described in Section 7.1.

Atom mappings are typically stored in the MetaCyc PGDB only, except for the reactions unique to other PGDBs.

3.7 Computation of Metabolite Gibbs Free Energies

The MetaCyc database provides the standard Gibbs free energy of formation for its compounds, and the change in Gibbs free energy for its reactions. These data were calculated by an algorithm within Pathway Tools. The algorithm first calculates the free energy of formation at pH 0 and ionic strength 0 ($\Delta_f G^0$) by using a technique based on the decomposition of the compounds into chemical groups with known free-energy contributions to the overall energy, based on the method of [26]. Then, the standard Gibbs free energy at pH 7.3 and ionic strength 0.25 ($\Delta_f G'^0$) is computed based on a technique developed by Robert A. Alberty [27]. In his technique, Alberty proposes to use several protonation states for some compounds, but we simplified the technique by always using only one protonation state, the state stored in MetaCyc. We use pH 7.3 because this is a common cellular pH, and computation of the protonation state of all compounds in MetaCyc were performed at that pH.

The change in standard Gibbs free energy of reactions, $\Delta_r G'^0$, is computed based on the $\Delta_f G'^0$ values of the compounds involved in the reaction. The $\Delta_f G'^0$ could not be computed

for some of the compounds in MetaCyc due to the impossibility of decomposing them into the groups provided by the technique of [26]. Consequently, the $\Delta_r G'^0$ is not computed for any reaction which has a substrate for which its $\Delta_f G'^0$ is not stored in MetaCyc.

3.8 Pathway/Genome Editors

The Editors support PGDB curation through interactive modification and updating of all the major data types supported by Pathway Tools. The editing tools included in Pathway Tools are as follows:

- Gene Editor: Supports editing of gene name, synonyms, database links, and start and stop position within the sequence.
- Protein Editor: Supports editing of protein attributes as well as of protein subunit structure and protein complexes (see Supplementary Figure 12). Enables users to assign terms from the Gene Ontology controlled vocabulary. Pathway Tools can store, edit, and display features of interest on a protein; see Section 4.9 for more details. When editing a protein feature the user selects a feature type (e.g., phosphorylation site), defines the location of the feature on the sequence, a bound or attached moiety where appropriate, a textual label, an optional comment, citations, and sequence motif.
- Reaction Editor: Supports editing of metabolic reactions, transport reactions, and signaling reactions. This editor checks reactions for elemental balance and charge balance.
- Pathway Editor: Enables users to interactively construct and edit a metabolic pathway from its component reactions. (See Supplementary Figure 11.)
- Signaling Pathway Editor: Enables users to interactively construct and edit a signaling-pathway diagram by using a toolkit of icons and operations inspired by CellDesigner [28] (See Supplementary Figure 14). Updates to the visual representation are automatically translated back to changes to component reactions and proteins.
- Regulation Editor: Enables definition of regulatory interactions including regulation of gene expression by control of transcription initiation, attenuation, and control of translation by proteins and small RNAs (see Supplementary Figure 13). This editor also allows creation of operons and definition of their member genes, as well as specifying the positions of promoters and transcription-factor binding sites.
- Compound Editor: Supports editing of compound names, citations, and database links. Pathway Tools has been interfaced to an external chemical structure editor, Marvin [29], which is a JAVA applet. A chemical compound duplicate checker runs whenever

chemical structures are entered or modified, to inform the user if the resulting structure duplicates another compound in that user's PGDB or in MetaCyc. Additionally, we have added support for displaying glycan structures in an icon-based style that follows the conventions of CFG (Consortium for Functional Glycomics). To edit these structures, we interfaced to another JAVA applet, GlycanBuilder [30, 31], to which we have added some extra functionality for better integration with Pathway Tools.

- Publication Editor: Supports entry of bibliographic references.
- Organism Editor: Supports editing information about the organism described by a PGDB, including species name, strain name, and synonyms; taxonomic rank within the NCBI Taxonomy; sample collection data (e.g., date, geographic location, host, body site); and phenotypic information, such as pathogenicity and relationship to oxygen.
- Cellular Architecture Editor: Enables users to specify exactly which set of cellular components are present in an organism or cell type, with appropriate defaults derived from the organism's taxonomy.
- Sequence Editor: Supports interactive, visual editing of the nucleotide sequence for a replicon, allowing insertion, deletion, and replacement of arbitrary sections of sequence. Coordinates of all objects affected by the edits are updated automatically.

3.8.1 Author-Crediting System

Often, multiple curators collaborate on development of a given PGDB, and it is desirable to attribute their contributions accordingly, both to identify who to ask if questions about particular entries arise, and to provide an incentive for high-quality contributions, because contributors will be able to clearly demonstrate their accomplishments.

Most Pathway Tools editors thus create *credits* of several kinds. When an object such as a pathway is first created, by default, a "created" credit is attached to the object, along with a timestamp. A credit for an object can refer to curators, to organizations, or to both. Pathway Tools provides a generated web page for every curator and organization that lists all the objects for which they are credited.

Other kinds of credits are "revised" when a curator substantially edits an object that was created some time ago, and a "last curated" flag can be set to indicate when a curator has last researched the literature available for a given object. The last-curated flag is useful for those objects about which almost nothing is known, to distinguish between the case where no curator ever looked at the object versus where an extensive search was performed but still nothing new was found. The "reviewed" credit is used to attribute reviews of DB objects by external domain experts.

3.9 Incorporation of Genome-Annotation Revisions

Some groups choose to store the authoritative version of the organism's genome annotation in a database external to the PGDB. Such users need the ability to incorporate revisions to the genome annotation into their PGDB without overwriting or otherwise losing any manual curation they added to the PGDB. Pathway Tools provides an interface for incorporating annotation revisions that takes as input one or more update files, either in GenBank format or as a PathoLogic Format file. The files can contain either a complete revised annotation for the organism, or they can contain just the information that has changed. The software will parse the update files and determine all differences between the new data and the old. Types of changes that are detected include new genes, as well as updated gene positions, names, synonyms, comments, links to external databases, and updated functional assignments. A graphical interface summarizes different classes of changes to the user, and gives the user the option of either accepting all updates (e.g., creating database objects for all the new genes) or of accepting/rejecting individual updates. Once this phase is complete and any changes to functional assignments have been made, the software will re-run the pathway-inference procedure described in Section 3.2, identify any new pathways that are inferred to be present and any existing pathways that no longer have sufficient evidence, and allow the curator to review those changes.

3.10 Consistency Checker and Aggregate Statistics

Pathway Tools contains an extensive set of programs for performing consistency checking of a PGDB to detect structural defects that sometimes arise within PGDBs. Also included in this component are tools for computing and caching aggregate statistics for a PGDB, such as computing the molecular weights of all proteins from their amino acid sequences. The statistics are cached so that they can be displayed quickly.

Roughly half of the programs automatically repair PGDB problems that they find. Such problems could be caused by either user data-entry errors or errors in Pathway Tools itself. Example checks include: ensuring that inverse relationship links are set properly (e.g., that a gene is linked to its gene product, and that the product links back to the gene); making sure pathways do not contain duplicate reactions; validating and updating GO term assignments with respect to the latest version of GO; performing formatting checks in comment text; searching gene reading frames for internal stop codons; and removing redundant bonds from chemical structures.

3.11 Schema Upgrading and Propagation of MetaCyc Updates

Most new releases of Pathway Tools include additions or modifications to the Pathway Tools schema that are made to model the underlying biology more accurately (such as adding

support for introns and exons), and to extend the data types within Pathway Tools (such as adding support for features on protein sequences). Because each new version of the software depends on finding data within the fields defined by the associated version of the schema, existing user PGDBs created by older versions of the software will be incompatible with these new software versions.

Therefore, every release of Pathway Tools contains a program to upgrade PGDBs whose schema corresponds to the previous version of the software, to the new schema version. For users who have not upgraded the software for several releases, several upgrade operations will be performed consecutively. Example upgrade operations include adding new classes to the PGDB from the MetaCyc PGDB; adding new slots to PGDB classes; deleting PGDB classes; moving data values from one slot to another; and moving objects from one class to another. The schema upgrade leaves the user's curated data intact.

Every new release of Pathway Tools includes a new version of the MetaCyc database, which, in addition to providing new data content, typically contains updates and corrections to existing pathways, reactions, and compounds. Pathway Tools includes an option to propagate such updates and corrections to an existing organism PGDB. However, because we do not want to override any manual edits made to a PGDB, this tool does not run automatically. Much like the tool for incorporating a revised genome annotation, described in Section 3.9, this tool organizes the changes into logical groups (such as all compounds with newly added structures, or all reactions with changed reaction equations), and allows the user to either accept an entire group of changes, or to examine and confirm individual members of a group.

4 The Pathway Tools Schema

The Pathway Tools schema defines structured representations of a broad range of biological data types to enable computational analyses and integration of many types of data. The schema consists of a set of classes and a set of slots. Classes describe types of biological entities, such as genes and pathways, and are arranged in a class–subclass hierarchy. Slots define attributes of PGDB objects and relationships between PGDB objects. Figure 2 provides an overview of the relationships among PGDB classes. For example, user queries can follow the relationship from a gene to the protein that it codes for, from a protein to a reaction that it catalyzes, and from a reaction to a metabolic pathway in which it is a component, to answer questions such as "find all metabolic pathways in which the products of a given gene play a role".

Every PGDB object has a stable unique identifier (ID)—a symbol that uniquely identifies that object within the PGDB. Example unique IDs include TRP (an identifier for a metabolite), RXN0-2382 (an identifier for a reaction), and PWY0-1280 (an identifier for a pathway). Relationships within a PGDB are implemented by storing object IDs within slots. For example, to state that the TRP (L-tryptophan) object is a reactant in the reaction RXN0-2382, a

slot of RXN0-2382 called LEFT (meaning reactants) contains the value TRP. Many PGDB relationships exist in both forward and backward directions (for example, the TRP object contains a field called APPEARS-IN-LEFT-SIDE-OF that lists all reactions in which TRP is a reactant). The slots LEFT and APPEARS-IN-LEFT-SIDE-OF are called inverses.

4.1 Metabolites, Reactions, and Pathways

There are two alternative ways in which one might choose to represent the metabolic network in a computer: as a simple listing of all metabolic reactions that occur in the cell, or by partitioning the reaction list into a carefully delineated set of metabolic pathways that describe small, functionally linked subsets of reactions. Which approach is preferred? Both approaches have value, and they are not mutually exclusive; therefore, Pathway Tools supports both views of metabolism in a PGDB.

Pathway Tools conceptualizes the metabolic network in three layers. The first layer consists of the small-molecule substrates upon which metabolism operates. The second layer consists of the reactions that interconvert the small-molecule metabolites. The third layer consists of the metabolic pathways whose components are the metabolic reactions of the second layer. Note that not all reactions in the second layer are included in pathways in the third layer, because some metabolic reactions have not been assigned to any metabolic pathway by biologists. Scientists who choose to view the metabolic network within a PGDB solely as a reaction list can operate on the second layer directly without interference from the third layer. The compounds, reactions, and pathways in levels 1–3 are each represented as distinct database objects within a PGDB. The relationships among the metabolic data types in a PGDB are depicted by the blue region of Figure 2.

The representation of reactions can capture atom-mapping information that records, for each atom in a reactant compound, its terminus atom in a product compound. The representation of metabolites can capture stereochemical structural information as well as glycan structures.

The pathways in PGDBs are modules of the metabolic network of a single organism. Pathway boundaries are defined by considering the following factors. Pathways are often regulated as a unit (based on substrate-level regulation of key enzymes, on regulation of gene expression, and on other types of regulation). Pathway boundaries are often defined at high-connectivity, stable metabolites [32]. Pathway conservation across multiple species is also considered, as are pathway definitions from the experimental literature [9].

4.2 The Proteome and the Genome

PGDBs define the proteome and the genome of an organism in the following manner, as depicted by the green region of Figure 2. The proteome of the organism is described as a set of PGDB objects, one for each gene product in the organism, and one for each complex

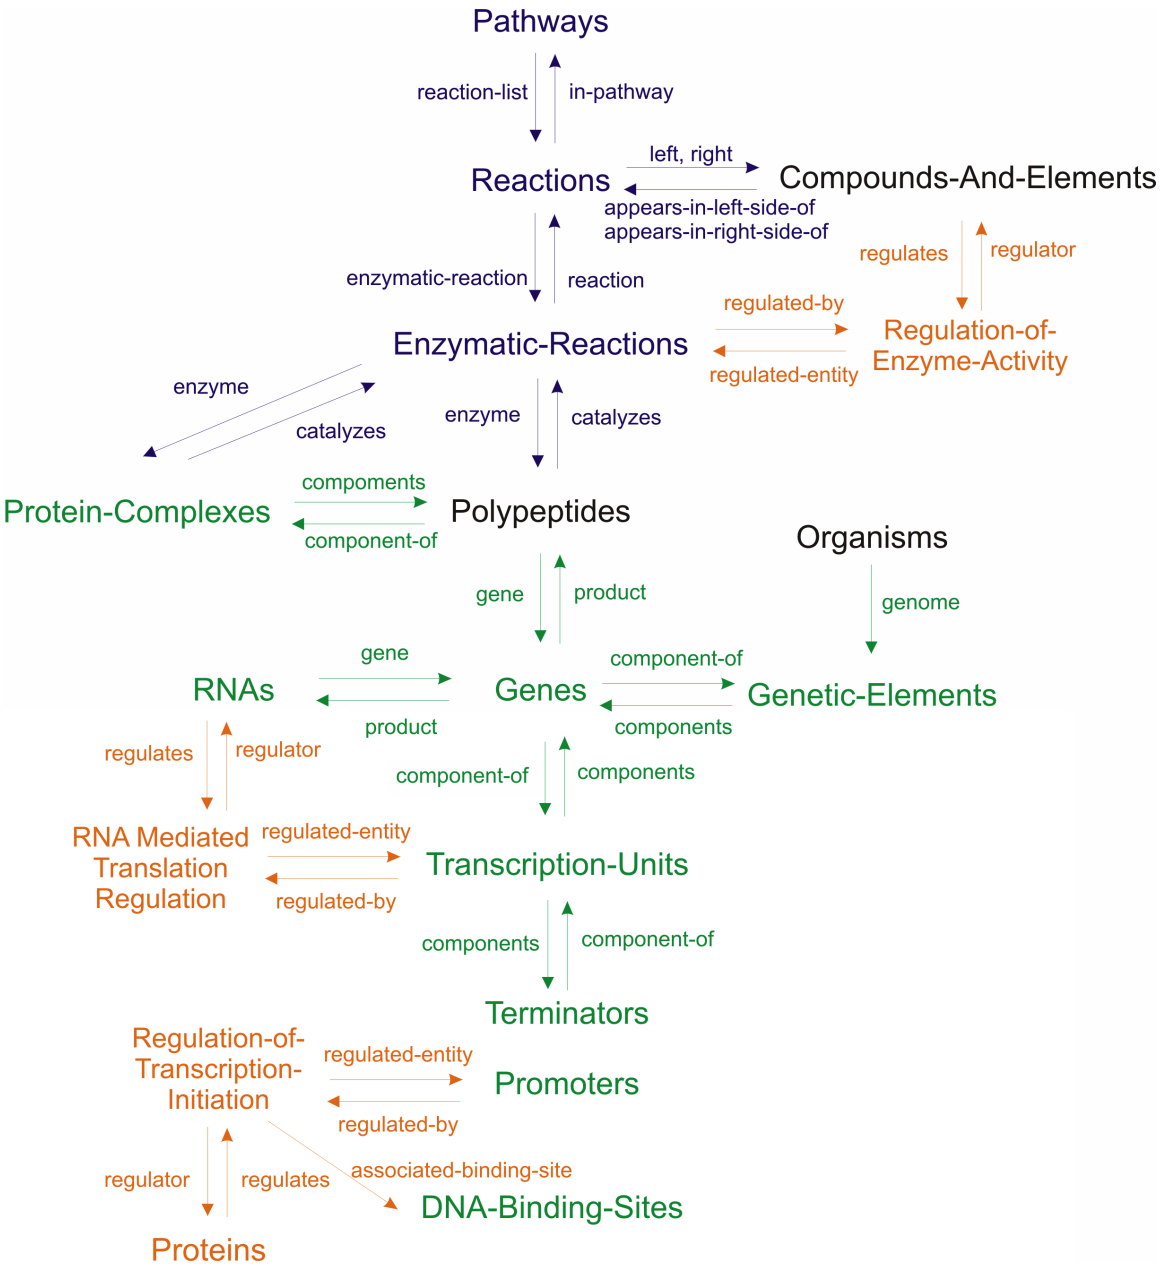

Figure 2: Major relationships among the major classes of the Pathway Tools schema. Colors indicate biological areas: blue for reaction and pathway information, green for genome and protein information, orange for regulation.

formed from two or more (identical or nonidentical) polypeptides. Furthermore, every chemically modified form of a monomer or of a multimer is encoded by a distinct PGDB object. For example, we might create one object representing an unmodified protein and another representing the phosphorylated form. Each protein object is in turn linked, through a slot in the object, to the metabolic reactions that it catalyzes. Proteins can also be substrates of reactions. Additional PGDB objects define features on proteins, as described in Section 4.9. Each protein product resulting from alternatively spliced forms of a gene is also represented by a distinct protein object. Each protein object records the exons of the gene that encodes it.

Protein objects are linked to gene objects that define the gene encoding each protein. Each gene in the genome is defined by a distinct PGDB object, as is every replicon (chromosome or plasmid) in the genome. Genes are linked to the replicon on which they reside. In addition, other features on the genome, such as operons, promoters, and transcription-factor binding sites, are described by PGDB objects.

The associations between enzymes and the reactions they catalyze are implemented by using an intermediary object called an enzymatic reaction, as shown in Figure 2. This arrangement enables us to capture the many-to-many relationship that exists between enzymes and reactions—one reaction can be catalyzed by multiple enzymes, and multifunctional enzymes catalyze multiple reactions. The purpose of the enzymatic reaction is to encode information that is specific to the pairing of the enzyme with the reaction, such as cofactors, activators, and inhibitors. Consider a bifunctional enzyme with two active sites, where one of the active sites is inhibited by pyruvate, and the second active site is inhibited by lactate. We would represent this situation with two enzymatic reactions linking the enzyme to the two reactions it catalyzes, and each enzymatic reaction would specify a different inhibitor.

4.3 Pathway Tools Regulation Ontology

The Pathway Tools schema can represent a wide range of regulatory interaction types. A regulation object within a PGDB captures information about each type of regulatory interaction. The available regulation types are as follows: (a) Substrate-level regulation of enzyme catalytic activity, such as the allosteric activation or competitive inhibition of an enzyme by a small molecule. Slots of this class identify the regulator molecule, the regulated enzymatic reaction object, encode the polarity of regulation (activation or inhibition), and the mechanism of regulation (allosteric, competitive, or noncompetitive). (b) Regulation of a bacterial promoter by a transcription-factor protein. The slots of this regulation class describe the transcription factor, the promoter that is regulated, and the binding site to which the regulator binds. (c) Regulation via premature termination of transcription (attenuation). This class of regulation is divided into six subclasses, each describing a different attenuation mechanism (e.g., ribosome-mediated, protein-mediated, or RNA-mediated). The slots of these classes identify the regulated terminator region, the regulator (a protein, RNA or

small-molecule, depending on the type of attenuation), and the regulator binding site if one exists. (d) Regulation of the translation of an mRNA transcript to the corresponding protein. This regulation class is divided into two subclasses to distinguish between regulation by a protein and regulation by a small RNA. The slots of these classes identify the regulated transcription-unit (which corresponds to a single transcript), the regulator protein or RNA, and the mRNA binding site where the regulator binds. An additional slot indicates whether regulation is by direct interference with the translation machinery, by processing of the mRNA transcript to promote or inhibit its degradation before translation, or both. (e) Regulation of protein activity by chemical modification, such as by phosphorylation, is represented by a reaction that converts the unmodified form of the protein to the modified form.

4.4 Conditions of Cellular Growth

A recent extension to the schema supports representation of conditions of cellular growth that include the chemical composition of the growth medium, pH, temperature, and aerobicity. This representation enables us to capture low-throughput information about conditions of cellular growth, and high-throughput information such as Phenotype Microarray [33, 34] datasets.

4.5 Gene Essentiality

A recent extension to the schema supports representation of gene-essentiality experiments. Our representation links growth phenotype (no growth, limited growth, or growth) under a given gene knockout with the conditions of cellular growth expressed as per Section 4.4. Multiple phenotypic observations can be recorded for a given gene knockout and growth condition to express conflicting experimental outcomes.

4.6 Organism Phenotype Data and Genome Metadata

A recent extension to the schema supports representation of microbial phenotypic data to enable users to query among the many genomes stored within a Pathway Tools website to find organisms pertinent to their research. Our representation adapts the MIGS [35] standard to incorporate metadata about the sample from which the organism was derived (e.g., geographic location, depth, health-or-disease state of host, human microbiome site), plus phenotypic information about the organism itself (e.g., relationship to oxygen, temperature range, and pathogenicity).

4.7 Pathway Tools Evidence Ontology

Database users want to know the type(s) of evidence that support assertions within a DB, and they want to know the strength of that evidence. We have developed an evidence ontology [36] that can encode information about *why* we believe certain assertions in a PGDB, the *sources* of those assertions, and the *degree of confidence* scientists hold in those assertions (although in practice the latter field is rarely populated). An example assertion is the existence of a gene in a PGDB. Was the gene predicted by using computational gene finding? Is it supported by wet-lab experiments? The Pathway Tools evidence ontology builds upon and substantially extends the Gene Ontology evidence ontology, which applies only to gene products.

Evidence about object existence in PGDBs is recorded as a structured *evidence tuple*. An evidence tuple enables us to associate several types of information within one piece of evidence. Each *evidence tuple* is of the form

```
Evidence-code : Citation : Curator : Timestamp : Probability
```

where `Evidence-code` is a unique ID for the type of evidence, within a hierarchy of 54 evidence types described in [37, 36]. `Citation` is an optional citation identifier such as a PubMed ID that indicates the source of the evidence. For computational evidence, the citation refers to an article describing the algorithm used. `Curator` identifies the curator who created this evidence tuple. `Timestamp` encodes when this evidence tuple was created. `Probability` is an optional real number indicating the probability that the assertion supported by this evidence is correct, such as a probability provided by an algorithm.

The Pathway Tools editors enable users to manually enter evidence codes, and the PathoLogic pathway and operon predictors annotate objects that they create with appropriate computational evidence codes. The Navigator supports display and querying of evidence codes.

4.8 Pathway Tools Cell Component Ontology

The Cell Component Ontology (CCO) is a controlled vocabulary of terms describing cellular components and compartments, and relationships between these terms [38]. It was developed to provide a controlled vocabulary of terms for annotating the subcellular locations of enzymes, and the compartments involved in transport reactions, in PGDBs. CCO spans all domains of life, and includes terms such as cytoplasm, cell wall, and chloroplast. The ontology currently contains 170 terms. CCO includes many terms and their definitions from the Gene Ontology [39], but substantially extends Gene Ontology. A recent extension to CCO enables any metabolic reaction to be annotated to one or more CCO compartments, to allow metabolic models to span multiple compartments.

4.9 Pathway Tools Protein Feature Ontology

We have developed an ontology of protein features to identify and represent post-translational modifications, binding sites, active sites, conserved regions, and other regions of interest on a protein. Starting from the list of feature types described in the UniProt User Manual [40], with some suggested additions from the SRI EcoCyc and MetaCyc database curators, we created an ontology of 40 feature classes.

Features fall into two major classes. For amino acid site features, the feature location is a list of one or more amino acid residue numbers (or residue types, if the feature is associated with a generic protein whose precise sequence is unspecified). For protein segment features, the feature location is a range defined by its starting and ending residue numbers.

Feature types that are classified as binding features (either covalent or noncovalent) permit specification of an attached group. The attached group could be a compound or compound fragment, as in the case of a protein that binds a small molecule. The attached group can also be another protein feature, as in the case of a disulfide bond or other cross-link between two features on different proteins, or any other type of molecule or binding site (such as a DNA binding site).

A different protein object is created in a PGDB for each biologically relevant modified form of a protein, and a single feature may be linked to multiple forms of the same protein. Some feature types are capable of existing in multiple states. For example, an amino acid modification feature can be in either the modified or the unmodified state (as in the case of a phosphorylation feature, which will be in the modified state when associated with the phosphorylated protein and the unmodified state when associated with the unphosphorylated protein), and a binding feature can be in either the bound or unbound state (as in the case of a metal-binding feature whose state indicates whether or not the metal ion is bound to the protein). We consider the state to be not an attribute of the feature, but rather an attribute of the pairing between a particular form of a protein and the feature.

5 Visualization and Querying of PGDBs

The Pathway/Genome Navigator component of Pathway Tools provides mechanisms for interrogating PGDBs and for visualizing the results of those queries. We begin by describing the query tools. We then describe visualization tools for individual biological entities (such as genes and pathways). We next describe the new SmartTables system, which enables the user to manipulate groups of PGDB objects, such as gene sets and metabolite sets. Finally, we describe system-level visualization tools that graphically display the entire metabolic network, entire regulatory network, and entire genome map of an organism.

The Navigator runs as both a desktop application and a web server, whereas the Pathway

Tools components described in some other sections run as desktop applications only (e.g., PathoLogic does not run as a web server). The desktop mode is faster, and has more overall functionality (see [41] for details), but the web mode has some functionality that is not present in the desktop mode.

5.1 Query Tools

Pathway Tools provides a three-tiered query paradigm, meaning that three different types of query tools are available, each of which represents a different trade-off between ease of use and query power. For example, the quick search is designed to provide a fast and simple way for new or casual users to find general information in the site.

The "Quick Search" box that appears at the top of most web pages generated by a Pathway Tools server is extremely easy to use. The user enters a search term and selects the organism whose PGDB the user wants to query. Pathway Tools searches that PGDB for objects whose primary name or synonyms contain the search term as a substring, and presents the list of results, organized by object type (e.g., gene, metabolite, pathway). The user can click on an object name to navigate to the display page for that object.

A set of intermediate-level query tools provides the ability to construct more powerful and precise searches against objects of a single class. One such query page exists for genes, proteins, and RNAs (see Supplementary Figure 1); additional query pages exist for pathways, reactions, chemical compounds, and growth media.

Finally, the Structured Advanced Query Page (SAQP) described in the next section enables advanced users to construct extremely powerful searches (that are approximately as powerful as those provided by the SQL language). The graphical interactive nature of this web form makes these searches much easier to construct than those using the SQL language.

5.1.1 Structured Advanced Query Page

The SAQP enables a biologist to construct precise structured searches. A query can be as simple as looking up a gene given a name, or as complex as searching several databases and several object types interconnected by several relations. The SAQP enables biologists to formulate queries whose power and expressiveness closely approach SQL, but without having to learn SQL. The SAQP translates a formulated query into BioVelo, an OQL-like language [42], before sending it to the web server.

The following explanation presents the elements of this web user interface using the example shown in Figure 3, involving a query against the class of protein monomers (Polypeptides) in the EcoCyc DB.

Step 1: Select Database and Class. The first step in building a query is to specify at

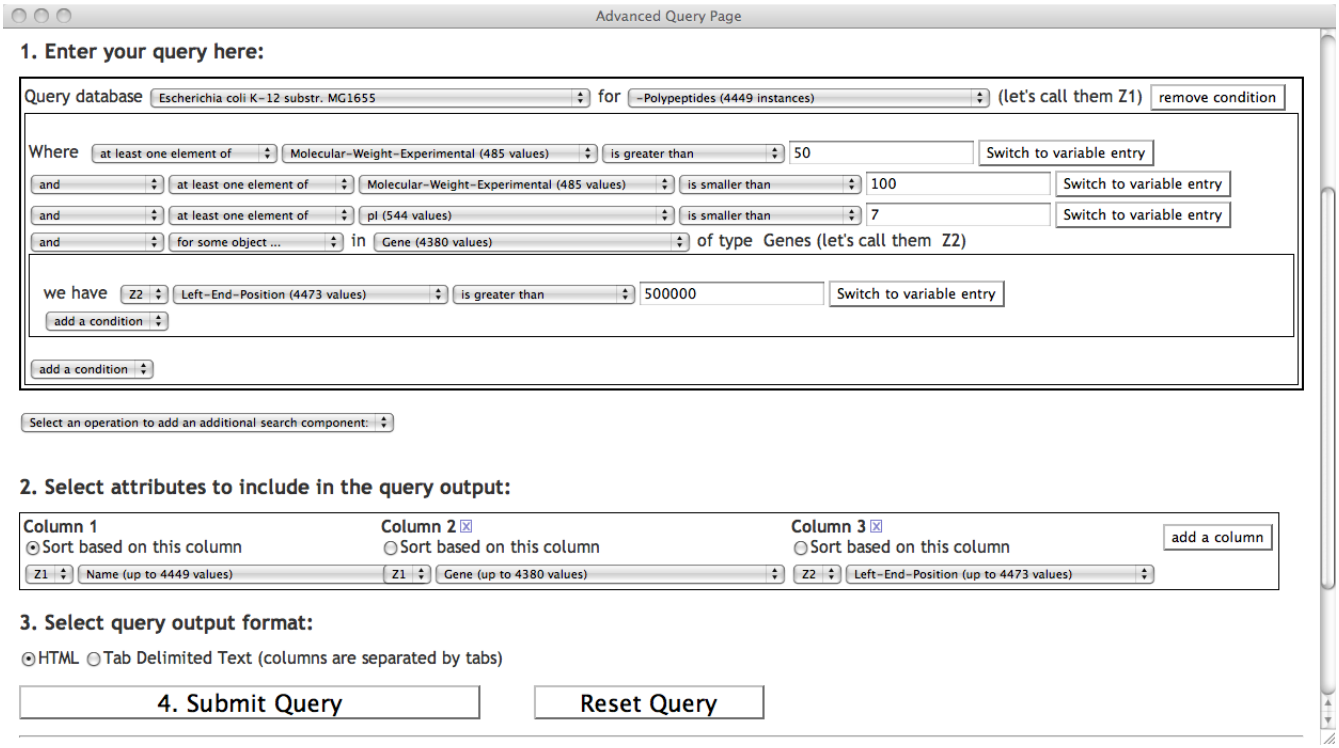

Figure 3: A query for the *E. coli* polypeptides whose experimental molecular weight lies between 50 and 100 kilodaltons, whose pI is smaller than 7, and whose gene is located after the first 500 kb of the genome. An output column is used to include the gene (or sometimes genes) producing each polypeptide using the second variable `Z2`.

least one database (DB) and the class of objects to search.

Step 2: Specify Conditions. Most queries include one or more conditions on the desired objects within the class. By clicking the button labeled `add a condition` in the initial blank SAQP, a *where* clause is added—visually boxed—in the search component. This operation adds a selector for an *attribute* (e.g., `name`) of the objects and a selector for a relational operator (e.g., `contains the substring`). It also adds a *free text box* to enter a number or string. Several other relational operators are provided, such as `is equal to`, and `is a substring of`. Regular expression matching is also available. This new field forms an *atomic condition*. Additional atomic conditions can be added to the query by using the button labeled "add a condition".

When selecting a relational operator, the list of relational operators provided is compatible with the type of the selected attribute. In the case of the attribute `name`, the selectable operators are for strings, because the *type* of the attribute `name` is string. The query in Figure 3 has three atomic conditions to filter the selected polypeptides.

Quantifiers on relations within the SAQP enable a join-like capability. For example, imagine that we want to extend the query with an additional restriction that depends on the *gene encoding the polypeptide,* not on the polypeptide itself. To do so, the user would add an

`and` condition and then select the `gene` attribute, which represents the gene encoding the polypeptide. We then select the quantifier operator `for some object...`, meaning that we want to define a condition that applies to some of the genes in the `gene` attribute of this polypeptide (although in the majority of cases only one gene will be present).

At this point, the SAQP adds a new indented query clause, to enable a condition to be defined on the gene. We have specified a constraint that its nucleotide coordinate must lie after the first 500 kb of the genome. Because several attributes and logical connectors can be specified in this new clause, forming a complex condition by itself, the web interface draws a box around this condition and introduces it with the `we have` keyword. A new unique variable, named `Z2`, is also introduced. This variable represents every value of the `gene` attribute.

Step 3: Define Query Results. The section titled `Select attributes to include in the query output` enables the user to describe the contents of the query results by selecting the attributes to display for each result object. The result of a query is a table containing zero or more rows, one for each query result. Each column in the table is a user-selected attribute.

5.1.2 Cross-Organism Search

We recently added a tool for searching across all organisms within a Pathway Tools web server, such as for the 5,500 organisms at BioCyc.org. The cross-organism search tool [43] searches for user-specified combinations of words in the Common-Name/Synonyms attributes, and/or the Summary attribute. It can search all types of objects in a given PGDB, or in user-specified object types, such as genes and/or pathways. It can search all organism databases present in the Pathway Tools web server, or it can search user-specified sets of organisms, such as all organisms within a selected taxonomic group. Indexing and searching is implemented using SOLR [44].

5.2 Visualization Tools for Individual Biological Entities

Genes/Proteins/RNAs: Since our last overview, we merged the individual display pages for a gene and its product (protein or RNA) into a single page that combines all information in one central place. These pages are quite extensive, listing information such as the map position of the gene on the chromosome, a graphical depiction of the chromosomal region containing the gene, and available gene-essentiality information. A new diagram, the regulation-summary diagram, integrates all known regulatory influences on the gene and gene product into a single figure. Common to all protein types is the ability to graphically display information about protein regions (such as phosphorylation sites and active sites) using a protein-feature ontology that we developed. **Enzymes:** The software displays the

reaction catalyzed by the enzyme and the name of the pathway that contains that reaction, if any (see Supplementary Figure 3); the activators, inhibitors, and cofactors required by the enzyme; and comments and citations for the enzyme. **Transporters:** The software displays the transport reaction catalyzed by the transporter (see Supplementary Figure 4). **Transcription factors:** The software displays diagrams for all operons controlled by the transcription factor (the regulon for the transcription factor) (see Supplementary Figure 5).

Reactions: Reaction display applies to metabolic, transport, and protein signaling reactions (see Supplementary Figure 2). The reaction display shows the one or more enzymes that catalyze the reaction, the gene(s) that code for the enzymes, and the pathway(s) that contain the reaction. The display shows the EC number for the reaction and the reaction equation. Conserved chemical moieties in the reaction substrates are drawn in different colors based on atom-mapping information.

Pathways: All metabolic pathway visualizations are computed automatically using pathway-layout algorithms. Pathway Tools can draw metabolic pathways at multiple levels of detail, ranging from a skeletal view of a pathway that depicts the compounds only at the periphery of the pathway and at internal branch points, to a detailed view that shows full structures for every compound, and EC numbers, enzyme names, and gene names at every reaction step (see Supplementary Figure 6). The user can customize a pathway drawing to include desired elements only, and to include superimposed omics data using "omics pop-ups" (see Figure 4). Support for display of signaling pathways (see Figure 5) was added recently; signaling pathway layout is performed by the user.

Electron Transfer Reactions and Pathways: A crucial role in cellular metabolism is played by electron transfer reactions (ETRs), which are of key importance in the energy household of a cell. In a series of redox steps, the high-energy electrons from some compounds drive the pumping of protons across a cell membrane, to maintain the proton motive force needed for ATP synthesis. Because electrons do not freely exist on their own, these redox reactions seem less intuitive than other small molecule reactions. We designed and implemented drawing code for a special ETR diagram, which shows the enzyme complex embedded in a membrane, and which schematically depicts the flow of electrons from one redox half reaction to another. Inside the membrane, the quinone/quinol cofactor is shown together with an indication of the cell compartments that are sources or sinks of the protons. An additional vectorial proton transport reaction can be added to the diagram. This results in displaying the flow of all substrates and products relative to the cellular compartments, in a similar way to what is customary in the biomedical literature. Pathways consisting of several ETRs joined together can also be depicted (see Figure 6).

Chemical compounds: The metabolite display shows the chemical structure for the compound (see Supplementary Figure 7). It lists all reactions in which the compound appears, and it lists enzymes whose activity is regulated by the compound.

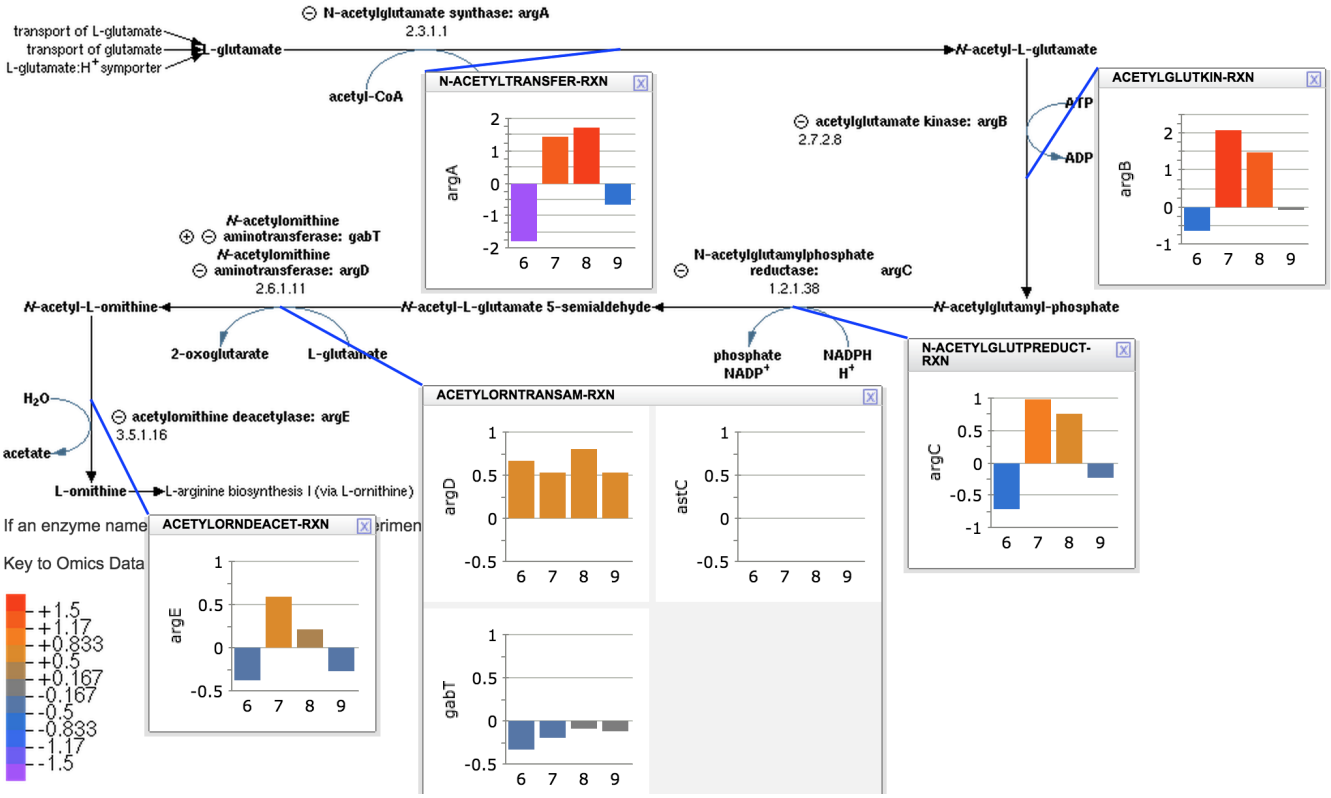

Figure 4: The EcoCyc L-ornithine biosynthesis pathway shown with omics pop-ups containing time-series data from a gene-expression experiment.

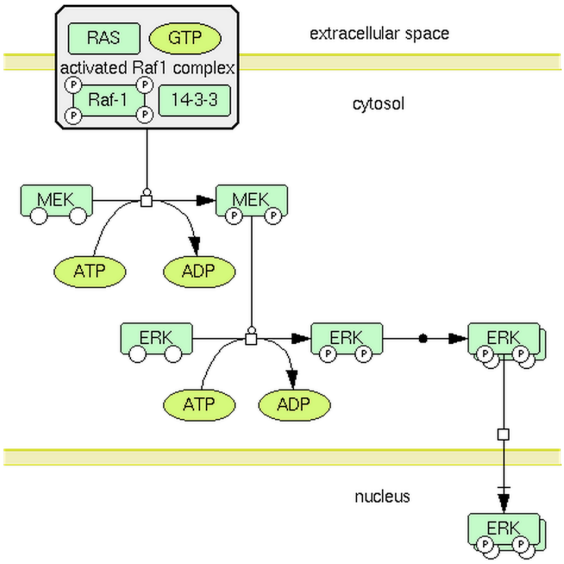

Figure 5: The HumanCyc MAP kinase cascade, a signaling pathway diagram.

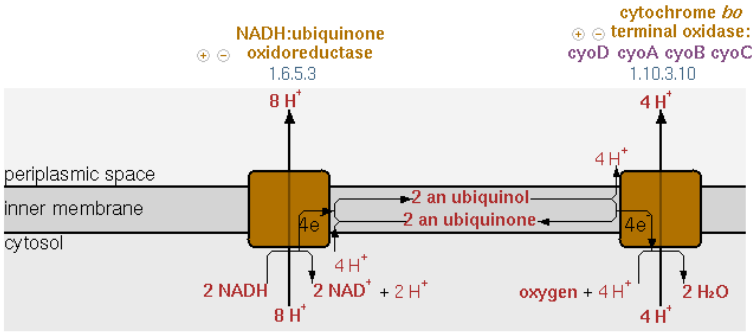

Figure 6: An electron transfer pathway diagram.

Transcription units: The display window for transcription units diagrams the transcription unit and its regulatory sites including promoters, transcription-factor binding sites, attenuators, and binding sites for proteins and RNAs that regulate its translation. The display contains sections describing each site within the transcription unit. The promoter section describes which sigma factor recognizes it. Sections for a transcription factor binding site describe its sequence, which transcription factor it binds, ligands that influence the activity of the transcription factor, and whether the effect of binding is to activate or inhibit transcription initiation. Sections for attenuators describe the signal that the attenuator senses, and show the sequence regions that form the attenuator.

5.3 SmartTables: Large-Scale Manipulation of PGDB Object Groups

SmartTables is a recent addition to Pathway Tools that enables users to construct and manipulate groups of PGDB objects through a spreadsheet-like user interface [45] (SmartTables were previously called Web Groups). SmartTables provide many powerful operations to biologist end users that previously would have required assistance from a programmer, and our user surveys indicated that SmartTables are reasonably easy for biologists to use [45]. Altogether, 2700 users of BioCyc.org have created more than 31,000 SmartTables.

A typical SmartTables use case is for a user to define a SmartTable by importing a list of PGDB objects from a file. For example, a user could define a metabolite SmartTable by importing a list of metabolites from a metabolomics experiment, where the metabolites are specified by metabolite name, BioCyc identifier, PubChem identifier, or KEGG identifier. (The set of objects in a SmartTable can also be defined from a query result, from any column of an existing SmartTable, or from the set of, say, all genes in a PGDB.)

The user can browse the set of objects in a SmartTable by paging through the table, and can modify the information displayed about each object by specifying which table columns to include (see Figure 7). SmartTable columns are derived from the PGDB attributes available

1 2 3 4 5 Next Show all

	All-Genes	Accession-1	Product	Regulation - direct regulators of gene
1	panZ	b3459	maturation factor for PanD	
2	ubiH	b2907	2-octaprenyl-6-methoxyphenol hydroxylase	RNA polymerase, sigma 70 (sigma D) factor
3	ubiG	b2232	UbiG [component of bifunctional 3-demethylubiquinone-*8* 3-*O*-methyltransferase and 2-octaprenyl-6-hydroxyphenol methylase]	CRP-cAMP DNA-binding transcriptional dual regulator RNA polymerase, sigma 70 (sigma D) factor
4	tyrA	b2600	TyrA [component of chorismate mutase / prephenate dehydrogenase]	TyrR-Tyrosine DNA-binding transcriptional repressor RNA polymerase, sigma 70 (sigma D) factor
5	trpE	b1264	anthranilate synthase component I	TrpR-Tryptophan DNA-binding transcriptional repressor RNA polymerase, sigma 70 (sigma D) factor an L-tryptophanyl-[tRNA^{trp}]
6	trpD	b1263	anthranilate synthase component II	TrpR-Tryptophan DNA-binding transcriptional repressor RNA polymerase, sigma 70 (sigma D) factor an L-tryptophanyl-[tRNA^{trp}]
7	trpC	b1262	indole-3-glycerol phosphate synthase / phosphoribosylanthranilate isomerase	TrpR-Tryptophan DNA-binding transcriptional repressor RNA polymerase, sigma 70 (sigma D) factor an L-tryptophanyl-[tRNA^{trp}]
8	trpB	b1261	tryptophan synthase, β subunit	TrpR-Tryptophan DNA-binding transcriptional repressor RNA polymerase, sigma 70 (sigma D) factor an L-tryptophanyl-[tRNA^{trp}]

Figure 7: A gene SmartTable. Column 1 shows the gene name, column 2 shows the *E. coli* genome "b-number" accession number for the gene (a property); column 3 shows the gene product name (a property). Column 4 shows the result of a transformation in which the regulator(s) of each gene were computed.

for each object, and can include information such as chemical structures, molecular weights, links to other databases, and nucleotide and protein sequence. A variety of filters and set manipulations are provided for SmartTables, such as removing or retaining all rows that match a user query; and computing the union, intersection, and set difference of two SmartTables. SmartTables are stored in the user's online web account, and a desktop version of SmartTables is also provided. SmartTables are private by default, but the user can make them public, share SmartTables with selected other users, or archive them in a frozen form in conjunction with a publication.

Several more advanced SmartTable operations are provided. *Transformations* compute new columns from relationships in a PGDB. For example, column 4 in Figure 7 is a transformation column that shows one or more regulators for each gene in column 1 that have been computed from PGDB relationships. Other gene transformations available include comput-

ing the metabolic pathways in which a gene's product occurs and computing the amino-acid changes caused by sequence variants. Different transformations are available for different datatypes. For example, the transformations available for a metabolite SmartTable include computing the reactions in which a metabolite occurs, the pathways in which a metabolite occurs, the proteins for which the metabolite is a ligand, and mapping the compounds to their equivalents in another PGDB.

A user can perform a statistical enrichment analysis on a gene or metabolite SmartTable to detect over-represented metabolic pathways or GO terms, or over-represented metabolic pathways, respectively. In addition, a SmartTable of genes or metabolites can be visualized on the cellular overview.

5.4 System-Level Visualization of Metabolic Networks

Pathway Tools can automatically generate organism-specific metabolic charts that we call Cellular Overview diagrams [2]. The diagram can be interrogated interactively and used to analyze omics datasets. Recently, the diagram was re-engineered for the web mode of Pathway Tools. The diagram can be generated as a PDF file for printing as a large-format poster. Example posters can be downloaded from [46].

Figure 8 depicts the Web Cellular Overview at low resolution painted with gene-expression data. It contains all known metabolic pathways and transporters of an organism (online example: [47]; example with animated display of omics data: [48]). Each node in the diagram represents a single metabolite, and each line represents a single bioreaction. Omics data (e.g., gene-expression or metabolomics measurements) for a given organism can be painted onto the cellular overview to place this data in a pathway context and to enable the user to discern the coordinated expression of entire pathways (such as the TCA cycle) or of important steps within a pathway. Omics pop-ups show all time points for particular reactions or metabolites of interest. Omics data may be loaded from a data file, GEO or PortEco dataset, or SmartTable, and superimposed on the Overview diagram for that organism. In web mode, the user has a choice of several color schemes—in desktop mode the color scheme is fully customizable.

Cellular Overview diagrams are generated automatically using an advanced layout algorithm [2]. Automated layout is essential to enable the diagram to accurately depict the underlying database content as that content evolves, without requiring time-consuming manual updates by curators that are bound to overlook some updates. In addition, automated layout enables generation of organism-specific cellular overviews that reflect the exact pathway content of each organism-specific PGDB in large PGDB collections such as BioCyc.

The Cellular Overview has many capabilities (described in more detail in [2]), including semantic zooming of the diagram (where the highest magnification corresponds to the detail shown in the poster version); highlighting of user-requested elements of the diagram (such as

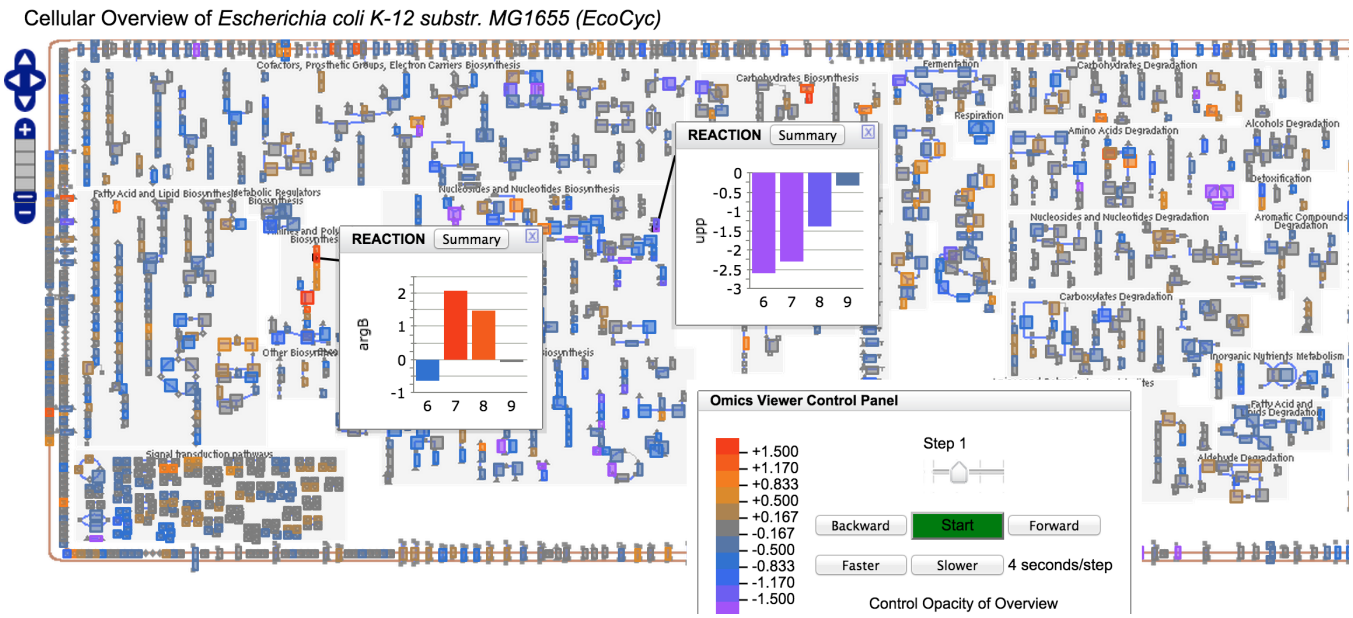

Figure 8: The Pathway Tools Cellular Overview diagram for EcoCyc, painted with gene-expression data. Omics pop-ups are shown for two genes.

metabolites or pathways); highlighting large, biologically relevant subnetworks (such as all reactions regulated by a given transcription factor); and highlighting comparative analysis results, such as comparison of the metabolic networks of two or more PGDBs. Because a metabolite can appear in several different places in the diagram, in desktop mode, the user can better visualize the flow of material through the metabolic network by selectively showing connections between a metabolite of interest, or all the metabolites in a pathway of interest, and everywhere else those metabolites appear in the network, as shown in Figure 9.

5.5 System-Level Visualization of Regulatory Networks

The Pathway Tools Regulatory Overview depicts the full transcriptional regulatory network stored in a PGDB in one screen, and enables the user to interrogate and explore relationships within the network. Figure 10 shows the Regulatory Overview for EcoCyc, after the user has asked the system to highlight all genes annotated under Gene Ontology term GO:0001539 (ciliary or flagellar motility). We can see that a few transcription factors control all *E. coli* motility genes.

5.6 System-Level Visualization of Genome Maps

The Pathway Tools genome browser displays a selected replicon (chromosome or plasmid), and enables the user to zoom into a region of the replicon by gene name or by coordinates.

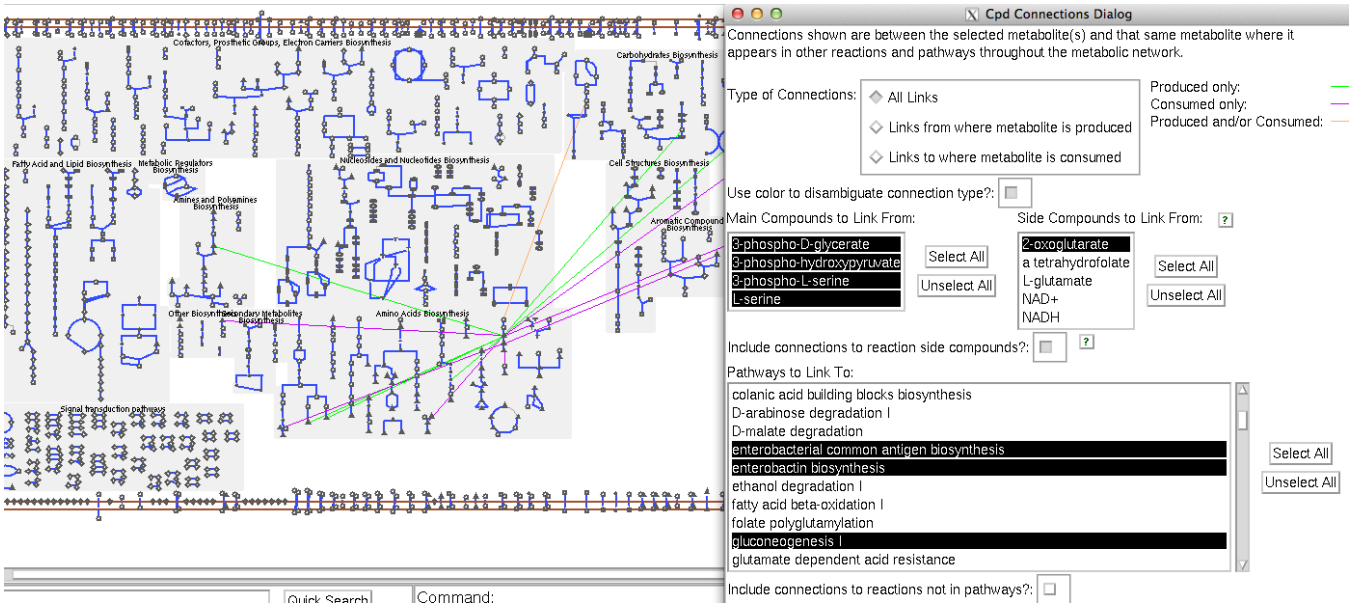

Figure 9: The Pathway Tools Cellular Overview diagram for EcoCyc, showing connections between metabolites of the serine biosynthesis pathway and selected other pathways.

The browser supports semantic zooming: as the user moves deeper into the genome, additional features are displayed, such as promoters and terminators. At high magnification, the genome sequence and the amino-acid sequence of coding regions become visible. Intron and exon boundaries are shown. The genome browser can be used in a comparative mode that displays replicon regions centered on orthologous genes across a user-specified set of genomes to show the genomic context of those genes (example: [49]). In comparative mode, the user retains the ability to navigate left or right in the genome, and to zoom in and out. The genome browser can also generate large-format genome posters in PDF format; see examples at [46].

The genome browser also supports display of tracks, meaning the ability to view positional data from external files along the genome, such as viewing predicted transcription-factor binding sites.

A user who zooms out far enough is presented with a depiction of all the genes on the replicon called the Genome Overview, shown in Figure 11. This diagram can be painted with omics data to provide a global genome view of large-scale datasets.

5.7 Sequence-Based Query and Visualization Tools

The following tools support query and visualization of sequence data from a Pathway Tools web server:

- Nucleotide Sequence Viewer: Gene pages include links to view or download the nu-

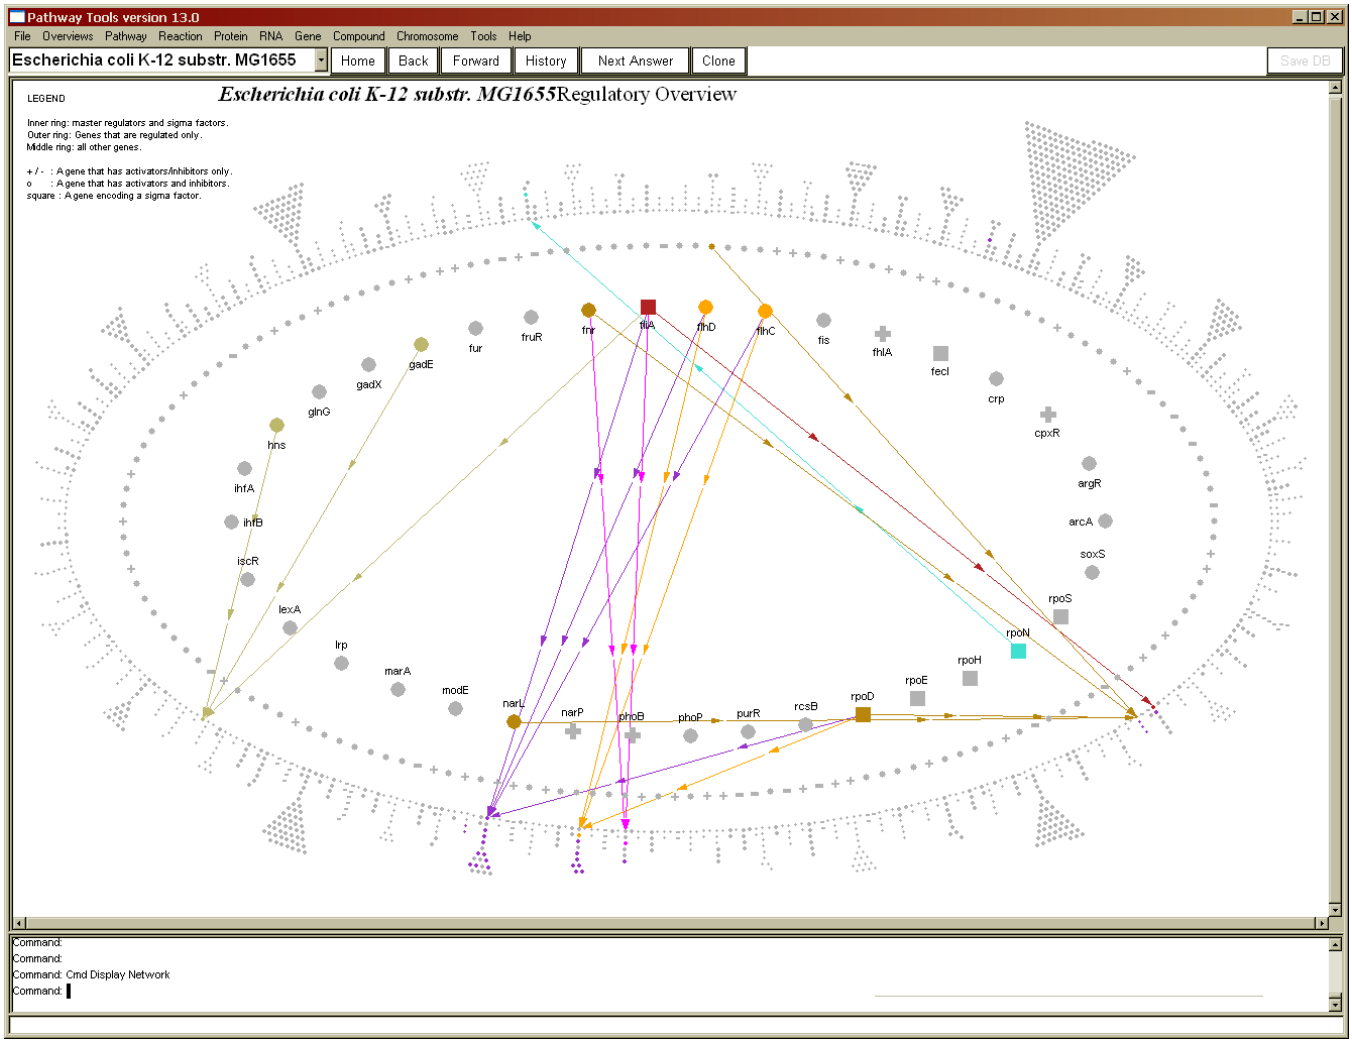

Figure 10: The Pathway Tools Regulatory Overview diagram for EcoCyc. The diagram depicts a full regulatory network as three concentric rings: the inner ring contains master regulator genes; the middle ring contains other regulators; the outer ring contains genes that are not regulators. An arrow (edge) from gene A to gene B indicates that gene A regulates gene B. Initially, no arrows are shown; the user can interactively add arrows, such as by clicking on a gene and requesting that arrows are added to genes that it regulates, or from the genes that regulate it.

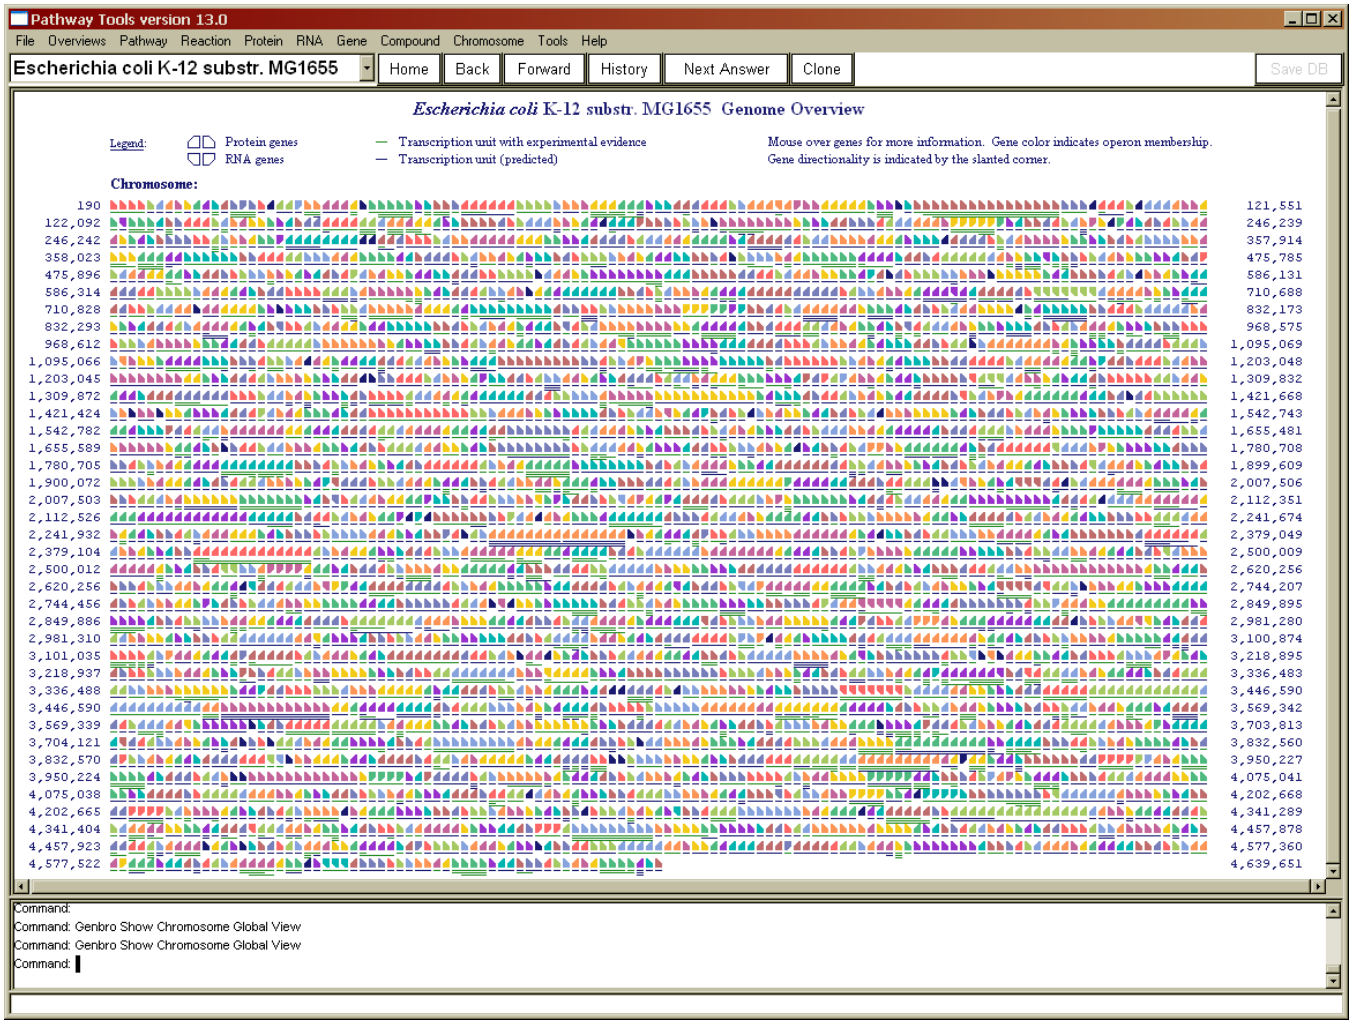

Figure 11: Pathway Tools Genome Overview diagram for EcoCyc. In this diagram genes are not shown to scale. Adjacent genes drawn in the same color are in the same operon. Left/right gene direction indicates transcription direction; up/down gene direction indicates genes coding for proteins versus RNAs. Horizontal lines under genes indicate transcript extents based on promoter and terminator information in the PGDB.

cleotide or RNA/protein sequence for the gene. When viewing the nucleotide sequence, an option is provided to include an additional upstream and/or downstream flanking region of any desired length. This option makes it easy to, for example, view the sequence of a regulatory region surrounding a gene of interest. Alternatively, the user can enter specific start and end coordinates and the desired strand to view the sequence of any arbitrary portion of the chromosome.

- BLAST Search: Users can use BLAST to search for all sequences within a single specified genome that match a query sequence.
- Sequence Pattern Search (PatMatch): The PatMatch facility enables searching within a single genome for all occurrences of a specified short nucleotide or peptide sequence (less than about 20 residues), with the ability to specify degenerate positions. The user can specify the kind and number of allowable mismatches, and whether to search coding regions only, intergenic regions, or the entire genome. Examples of situations in which this facility might be useful include searching for all occurrences of a particular regulatory motif upstream of any gene, or all occurrences of a known cofactor binding motif within proteins.
- Multiple Sequence Alignment: From a gene page, users can request a multiple sequence alignment between the nucleotide or amino acid sequences of that gene and its orthologs in a user-specified set of organisms. Alignments are displayed using MUSCLE [50].

6 Computational Access to PGDB Data

In addition to the user-friendly graphical interfaces to PGDBs provided through the web and desktop versions of Pathway Tools, the software supports multiple methods for importing and exporting data from files and via programmatic interactions, which are summarized in Figure 12.

Programmatic Access through Application Programming Interfaces (APIs). Programmers can access and update PGDB data directly [14, 51] by writing programs in the Python, R, Java, Perl, and Common Lisp languages. R, Java, and Perl queries are executed using systems called RCyc [52], JavaCyc [53], and PerlCyc [54].

PythonCyc is a Python package that enables programmatic access to Pathway Tools. The package provides the basic functions to access and modify the data of any PGDB. It also exposes more than 150 functions of Pathway Tools, among them, the MetaFlux module. The PythonCyc package is hosted on GitHub and is a separate installation from Pathway Tools. A full API documentation and a tutorial is available online. Please consult the URL [55] for access to the package, the API documentation and the tutorial.

Data Import Formats. Pathway Tools can import data from many sources into a PGDB.

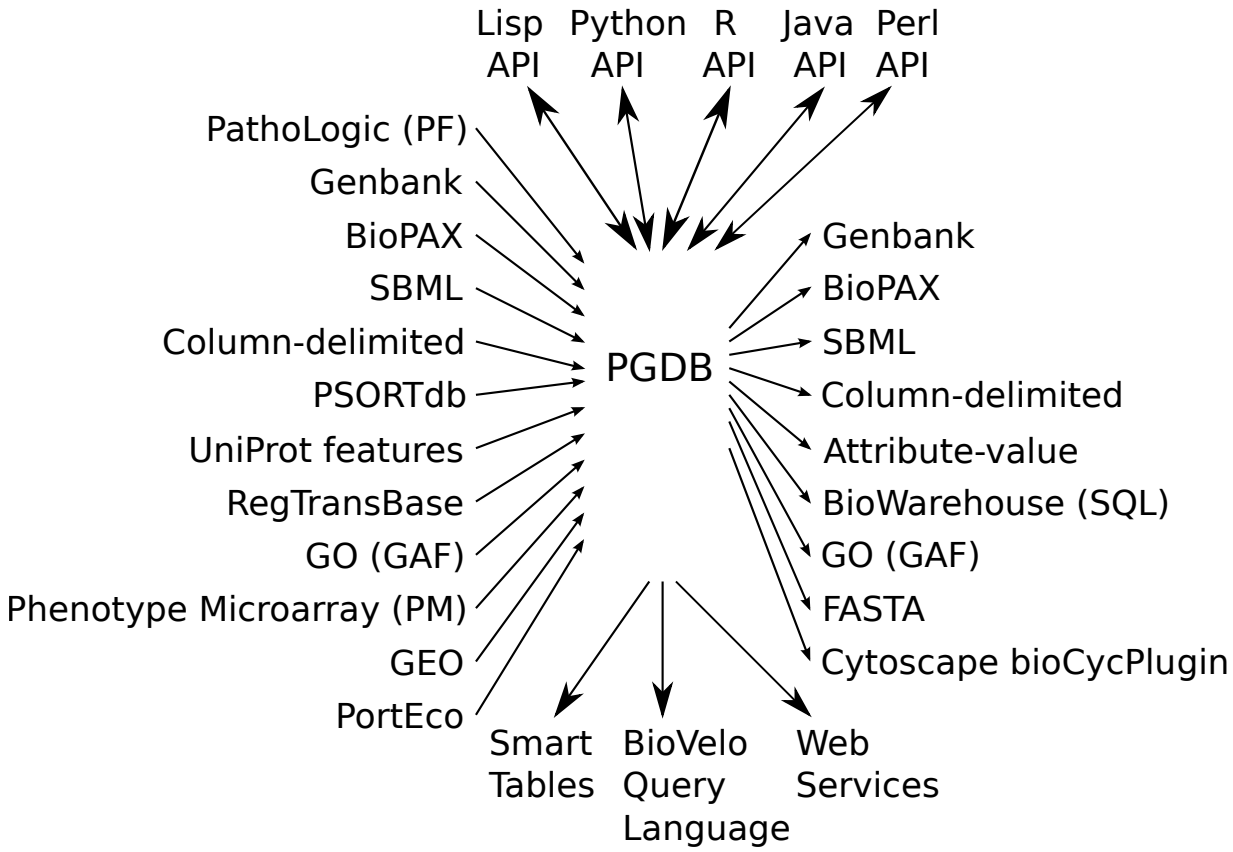

Figure 12: Pathway Tools supported formats and APIs for data import and export.

First of all, PathoLogic can create or update a PGDB based on a PathoLogic Format (PF) [56] or GenBank file [57]. A PGDB can also be populated from an SBML [58, 59] or BioPAX [60, 61] file. Bulk updating of various data items can be accomplished by a simple tab-delimited table format.

Several specialized data types are also supported. The following sources can be imported for proteins. UniProt sequence annotation features [62] can be fetched from a Biowarehouse [15] server that was loaded with SwissProt and TrEMBL. GO (Gene Ontology) annotations can be loaded from a GAF [63, 64] file. PSORTdb [65] cellular localization data can be imported from tab-delimited files.

Gene regulatory data can be imported from RegTransBase [66], which is a SQL database. Growth conditions versus growth media can be imported from Phenotype Microarray (PM) [33, 34] files. High-throughput expression data can be obtained from NCBI GEO [67] and PortEco [68] via web services.

File Export Formats. Pathway Tools can export PGDBs into several file formats that we have developed, which include tab-delimited tables and an attribute-value format (see [69]). Pathway Tools can also export subsets of PGDB data to other common formats including SBML [58, 59], BioPAX [60, 61], GO (in the GAF format) [63, 64], GenBank [57], and

FASTA [70]. The bioCycPlugin for Cytoscape [71] makes use of web services in conjunction with BioPAX, to select and export pathways into the Cytoscape environment.

Export for Relational Database Access via BioWarehouse. For scientists who want to query PGDB data through a relational database system, the attribute-value files exported by Pathway Tools can be loaded into SRI's BioWarehouse system [15]. BioWarehouse is an Oracle or MySQL-based system for integration of multiple public bioinformatics databases. PGDB data can be queried through BioWarehouse alone or in combination with other bioinformatics DBs such as UniProt, GenBank, NCBI Taxonomy, ENZYME, and KEGG.

Queries Using the BioVelo Query Language, and Web Services. Pathway Tools provides a powerful database query language for PGDBs, called BioVelo [72, 73]. BioVelo queries can be issued through an interactive web form, and through APIs.

Pathway Tools web services [74] enable programmatic retrieval of numerous data types, based on submitted HTTP GET or POST commands, and have expanded substantially in recent years. Users can access BioVelo queries, a Metabolite Translation Service, and can also invoke omics visualization services and SmartTable manipulations via web services. Services for SmartTables include creation, retrieval, copying, and deletion; applying many transformations; and changes like adding and deleting rows, columns and cells.

7 Graph-Based Metabolic Network Analyses

This section describes Pathway Tools modules for performing graph-based analyses of metabolic networks.

7.1 Metabolic Route Searching Using The RouteSearch Tool

RouteSearch [75] is a Pathway Tools component that enables the exploration of the reaction network of a PGDB, and engineering of new metabolic pathways. RouteSearch computes optimal metabolic routes (that is, an optimal series of biochemical reactions that connect start and goal compounds), given various cost parameters to control the optimality of the routes found. RouteSearch can display several of the best routes it finds using an interactive graphical web page (see Figure 13 for an example).

When RouteSearch is used in metabolic engineering, it uses the MetaCyc database as its external reaction database for new reactions to include in an organism. The cost for adding one reaction is selected by the user. Typically this cost, an integer, is larger than the cost of using one reaction from the organism. In such a case, a new reaction would be added for a route if it conserve more atoms from the start compound to the target compound. The cost of losing one atom from the start compound is also selected by the user, and it is

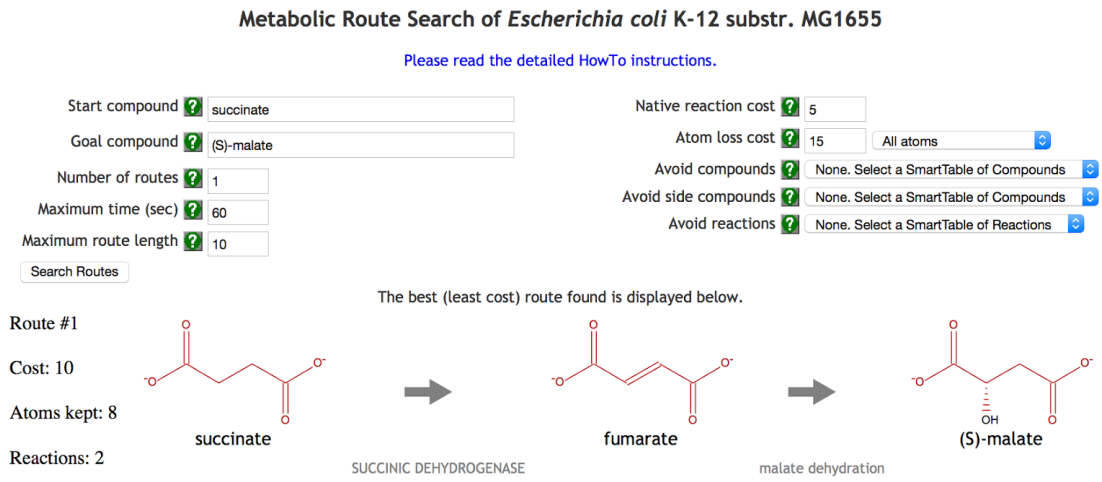

Figure 13: The RouteSearch web interface is shown with the result of one short pathway found. The arrows represent reactions and are tagged with the protein names catalyzing them. All atoms are conserved from the start compound (succinate) to the target compound ((S)-Malate) in this simplified example.

typically larger than the cost of one reaction from the organism or the external library of new reactions.

In computing optimality, RouteSearch takes into account the conservation of non-hydrogen atoms from the start compound to the goal compound. The more atoms that are conserved, the more efficient the transformation from start to goal. To compute the number of conserved atoms, RouteSearch uses pre-computed atom mappings of reactions that are available in MetaCyc. An atom mapping of a reaction gives a one-to-one correspondence of each non-hydrogen atom from reactants to products. RouteSearch is available only in web mode in Pathway Tools.

7.2 Dead-End Metabolite Analysis

Dead-end metabolites are metabolites that are only produced by the metabolic network or only consumed by the network. More precisely, the tool for identifying dead-end metabolites will report a small-molecule metabolite M as a dead-end metabolite in the cellular compartment C if and only if one of the following conditions is true:

1. M or parent classes of M are only **consumed** by small-molecule reactions occurring in C, and M or parent classes of M are not transported into C.

2. M or parent classes of M are only **produced** by small-molecule reactions occurring in C, and M or parent classes of M are not transported out of C, and no enzyme in the PGDB uses M as a cofactor

7.3 Computation of Blocked Reactions

The MetaFlux component of Pathway Tools computes the blocked reactions in a reaction network (see 10.6.1)—reactions that can never carry flux because of blockages in the network.

7.4 Prediction of Network Choke Points

One application of a metabolic network model is to find network bottlenecks, which if blocked could kill the cell. Such bottlenecks could constitute antimicrobial drug targets. We have developed a tool for predicting these so-called choke points.

The Pathway Tools choke-point detection algorithm examines the reactions attached to a given metabolite, and processes one metabolite at a time. The first step is to assemble the list of metabolites to examine. This is done by collecting (1) all reactions that are in pathways, plus (2) reactions that stand alone, but which use only small molecule metabolites. The reactions that came from pathways may use some macromolecular substrates, such as proteins that are modified by the reaction. From this list of reactions, the algorithm collects all of their substrate metabolites (meaning their reactant or product metabolites).

Definition [16]: A "choke point reaction" is a reaction that either uniquely consumes a specific substrate or uniquely produces a specific product in a metabolic network, and is also balanced by at least one reaction that respectively produces or consumes that substrate. Specifically, the algorithm searches for two types of choke point reactions: (a) Reactions R_1 such that only a single reaction R_1 produces metabolite M, and at least one reaction consumes M. (b) Reactions R_2 such that only a single reaction R_2 consumes metabolite M, and at least one reaction produces M. These definitions imply that to find a choke point, all reactions involving M must be unidirectional. These choke point reactions are collected and returned as the result. Note that the definition excludes reactions directly connected to dead-end metabolites.

The resulting candidate choke point reactions can be painted onto the cellular overview to facilitate further analysis.

8 Comparative Tools

Pathway Tools contains a rich set of operations for comparing the information in two or more PGDBs. These operations range from comparison of genome-related information to comparison of pathway information. These comparisons are of several types.

The comparative genome browser discussed in Section 5.6 displays replicon regions centered on orthologous genes across a set of genomes (see Supplementary Figure 9).

The user can generate a comparative table for a given metabolic pathway across a specified set of organisms. For each organism the table shows the presence of pathway enzymes and operon structures of genes within the pathway.

A global comparison of the metabolic networks of multiple PGDBs can be performed by highlighting on the Cellular Overview diagram (see Section 5.4). This tool enables the user to highlight in the Cellular Overview reactions that are shared, or not shared, among a specified set of organisms.

Finally, a general comparative analysis facility enables the user to generate comparative report tables for many aspects of a PGDB. As well as being used for comparative analyses, these tools can be used to generate statistics regarding the content of a single PGDB. These tools are general in that they present their results in a standard format, and they enable the user to drill down to specific results in a consistent fashion. The initial report page shows summary statistics, but the user can drill down to compare all instances of a category by clicking on elements of a report table.

For example, consider the transporter report page in Supplementary Figure 10. Table 2 within that report summarizes the number of uptake transporters found in two organisms. A user who wants to see the actual transported substrates clicks on the text "Compounds transported into the cell" to generate a new report page containing a table listing the union of all substrates imported by both organisms, along with an indication of which organisms transport each substrate, and which transporter is utilized. If the user clicks on a data cell within Table 2, such as the number of imported substrates in *E. coli* K-12 (156), a page is generated that lists those substrates only. Similar functionality applies to most tables in these reports.

The following report types are provided. An example comparative report is available at URL [76].

- Reaction report includes the following statistics for each selected organism:
  - Number of reactions containing substrates of different types (e.g., reactions for which all substrates are small molecules, and for which some substrate is a protein or a tRNA)
  - Number of reactions in each Enzyme Commission (EC) category
  - Number of reactions containing different numbers of isozymes
- Pathway report includes these statistics:
  - Number of pathways in each category within the MetaCyc pathway ontology
  - Number of pathways with different numbers of pathway holes
- Compound report includes these statistics:

  - Frequency with which different compounds appear in different metabolic roles (substrate, cofactor, inhibitor, activator)

- Protein report includes these statistics:
  - General statistics on number of monomers versus multimeric complexes, breakdown of multimers into heteromultimers and homomultimers
  - Statistics on multifunctional enzymes

- Transporter report includes these statistics:
  - Number of efflux versus influx transporters
  - Number of genes whose products are transporters
  - Number of unique transported substrates, both overall and broken down by efflux versus influx
  - Number of transported substrates that are substrates in metabolic pathways or are enzyme cofactors
  - Transporters with multiple substrates, and substrates with multiple transporters
  - Operon organization of transporters

- Ortholog report includes these statistics:
  - List of all orthologous proteins across the selected organisms
  - Proteins that are shared in all selected organisms, or unique to one organism

- Transcription Unit report includes these statistics:
  - Distribution of number of genes per transcription unit
  - Distribution of number of operons into which metabolic pathway genes are distributed

9 Software and Database Architecture

Pathway Tools is mostly implemented in the Common Lisp programming language,[2] with some components implemented in JavaScript. We chose Common Lisp because it is a high-productivity programming environment. Because Lisp is a very high-level language, one line of Lisp code is equivalent to several lines of code in a language such as Java or C++. Therefore, the same program can be written more quickly in Lisp, with fewer bugs. A study by Gat found that compiled Lisp programs generally run faster than Java programs, and

[2]We use the Allegro Common Lisp implementation from Franz Inc., Oakland, CA.

that a given program can be developed 2–7 times faster in Lisp than in Java [77]. Common Lisp also has a very powerful interactive debugging environment.

Lisp has powerful dynamic capabilities that are illustrated by a Pathway Tools feature called auto-patch. Imagine that a Pathway Tools user site has reported a bug in the software. Once our group has found a fix for the bug, we put a patch file that redefines the offending Lisp function(s) on the SRI website. The next time Pathway Tools is started at remote sites, it automatically downloads the patch (in compiled form) from the SRI website, puts the patch in an appropriate directory, and dynamically loads the patch file into the running Pathway Tools to redefine the altered function(s).

Pathway Tools consists of 640,000 lines of Common Lisp code, organized into 20 subsystems. In addition, 24,000 lines of JavaScript code are used within the Pathway Tools web interface. Pathway Tools runs on the Macintosh, Linux, and Microsoft Windows platforms. Pathway Tools was ported to 64-bit architectures in recent years.

The main bioinformatics modules of Pathway Tools are the Navigator, Editors, and Patho-Logic, plus a chemoinformatics subsystem that includes tools such as SMILES [78] generation and parsing, a chemical substructure matcher, plus a large set of shared utilities that we call the Pathway Tools core. Pathway Tools uses an object-oriented database system called Ocelot. The Pathway Tools user interface relies on a graph layout and display package called Grasper [79], and web and desktop graphics packages called CWEST and CLIM (the Common Lisp Interface Manager). Other software used by and included with Pathway Tools are (bioinformatics) textpresso; MUSCLE [50]; PatMatch [80]; BLAST; libSBML [81]; (chemoinformatics) Marvin [29]; GlycanBuilder [31]; InChI [82]; (lisp) ARNESI; 5am; cl-json; cl-store; (other) SCIP [83]; ghostscript; SKIPPY [84]; Yahoo User Interface library (YUI) [85]; SOLR [44]; and MySQL.

Ocelot is an object/relational database management system (DBMS) developed by our group at SRI [4, 18]. Ocelot combines the expressive power of frame knowledge representation systems [86] developed within the AI community (whose object data model is far superior to the relational data model for representing biological data[3]) with the scalability of relational database management systems (RDBMSs). Ocelot DBs are persistently stored within an

[3]Superior aspects of the object data model include the following. The object data model is better at managing very complex schemas. That is, if the same domain is represented within the object data and within the relational model, the object schema is usually much more compact and easier to comprehend. One reason is that inheritance enables the object data model to define subclasses by extending existing classes (e.g., the class Polypeptides is a subclass of the class Proteins), whereas the relational model would force attributes shared between the two tables to be duplicated in each, which both obscures the fact that the two tables are related, and complicates schema evolution. Relational normalization also increases the size of the schema by forcing the creation of new tables for every multi-valued attribute, which is not required in the object data model. The object data model used by Ocelot is particularly flexible in supporting any type of schema evolution without forcing the entire database to be reloaded (unlike relational DBMSs), which is important in bioinformatics because the complexity of biological data forces never-ending enhancements to the schema (note that not every object DBMS provides such flexibility).

Oracle or MySQL RDBMS. Ocelot objects are faulted on demand from the RDBMS, and in addition are faulted by a background process during idle time. Objects that were modified during a user session are tracked and saved to the RDBMS during a save operation. Ocelot uses optimistic concurrency control [18]—during a save operation it checks for conflicts between the updates made by the user and updates saved by other users since the saving user began their session or last made a save operation. This approach avoids the overhead of locking that becomes problematic in object databases because modifications to one object often cascade to related objects and could require a large number of lock operations. The optimistic concurrency control works well in practice because curators tend to focus in different biological areas and therefore rarely update the same objects at the same time.

Ocelot DBs can also be saved to disk files, in which case the RDBMS is not needed (see Figure 14). The file persistence configuration is simpler to use since it does not require purchase or installation of an RDBMS. It provides an easy and low-cost way to begin a PGDB project; a project can switch to an RDBMS configuration as its complexity grows. The advantage of an RDBMS configuration is that it provides Ocelot with multi-user update capabilities, and it permits incremental (and therefore faster) saving of DB updates. The RDBMS configuration also enables Ocelot to maintain a history of all DB transactions — DB curators can examine the history of all updates to a given object to determine when a given change was made, and by whom. This functionality is very useful when diagnosing mistakes within a PGDB.

Figure 15 shows the graphics architecture of Pathway Tools. The Grasper graph toolkit is used in pathway layouts, and in the cellular overview and regulatory overview. Grasper graphics, and all other graphics generated by Pathway Tools, are rendered using the CLIM Common Lisp graphics system, which is implemented using the X window system on Linux and Mac, and the native Windows API on Windows. When Pathway Tools runs as a desktop application, CLIM graphics directly update the user's screen.

Pathway Tools can also run as a web server, which is how it powers websites such as BioCyc.org. Pathway Tools does not run with an associated HTTP server such as Apache. Instead, Pathway Tools provides a fully functional web server that includes services such as compression and connection keep alive. HTTP servers typically start a new operating-system process for each incoming web request that terminates after the request has been serviced. In contrast, Pathway Tools starts one long-lived web server process that can service many thousands of web requests; that process forks an internal thread to service each incoming request.

Pathway Tools processes an incoming web request in the following manner. The top-level directory name within an incoming URL indicates whether the operation is requesting a static file or dynamically generated page.

- **Static files:** A small number of web pages, such as the home page and informational pages for BioCyc.org, are implemented as disk files. Pathway Tools can serve file-based

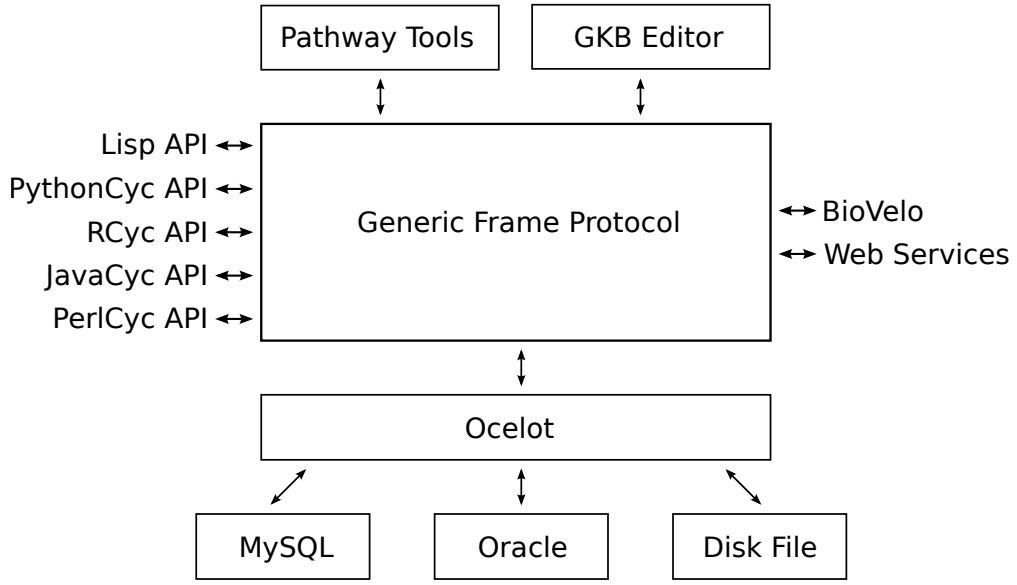

Figure 14: Database architecture and APIs for Pathway Tools.

web pages like a traditional web server.

- **Dynamically generated pages:** Most Pathway Tools pages are generated dynamically by processing that typically includes querying PGDBs within the Pathway Tools virtual memory and generating query outputs and visualizations, often using the same code as for desktop mode. The CLIM graphics generated by software such as the pathway layout code are dynamically converted to HTML and GIF images using CWEST [87], which uses Skippy to generate GIF images. The resulting HTML and GIF images are returned to the user's web browser, and the Pathway Tools web server awaits the next query. The GIF images include generated specifications of mouse-sensitive regions and of what operations should be invoked when the user clicks on such a region.

10 Metabolic Modeling with MetaFlux

The MetaFlux component of Pathway Tools is used to develop and execute quantitative metabolic flux models for individual organisms and for organism communities. MetaFlux uses the steady-state modeling technique of flux-balance analysis, which can be used to predict the phenotypes of an organism, or a community of organisms, based on a specification of available nutrients in the growth environment.

MetaFlux offers several modes of operation:

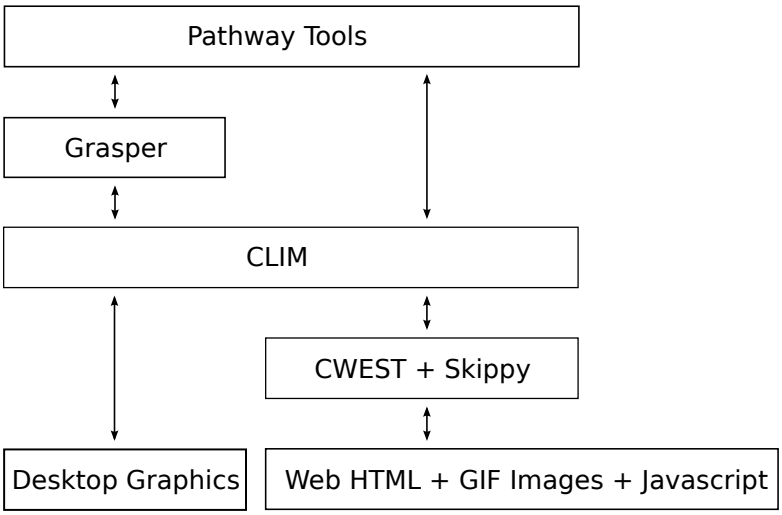

Figure 15: Graphics architecture of Pathway Tools.

1. Solving mode: execute a metabolic model for a single organism or for a community of organisms
2. Development mode: generate hypotheses on how to fill gaps in a developing metabolic model
3. Knockout mode: run metabolic models under gene-knockout scenarios

MetaFlux can be controlled via a graphical user interface and via a Python API.

The next section provides an overview of metabolic model development using MetaFlux.

10.1 The MetaFlux Model Development Process

Pathway Tools provides a unique environment for the development of metabolic flux models for several reasons. First, it includes a range of tools that support fast and accurate development of metabolic models from annotated genomes. Second, metabolic models developed with MetaFlux are highly accessible to the user, and are coupled with extensive enriching information, resulting in models that are easier to understand and reuse.

The high-level steps for developing a metabolic model from an annotated genome using Pathway Tools are as follows. For more information on the genome-scale metabolic reconstruction process, we suggest the comprehensive COBRA-based [88] overview published by Thiele and Palsson [89].

1. The PathoLogic tool computes a metabolic reconstruction by inferring the reactome of an organism from its annotated genome. By combining the enzyme-name matching

tool with the extensive reaction information in the MetaCyc database, we obtain a very complete mapping of annotated enzymes to metabolic reactions.

2. Subsequent inference of metabolic pathways by PathoLogic fills a significant number of missing (gap) reactions, because as well as inferring the presence of pathways, pathway inference also infers the presence of pathway reactions that were not initially identified.

3. The pathway hole filler identifies enzymes that catalyze those pathway-hole reactions. (Note this step is optional and is informative in nature, because it does not modify the set of reactions in the model.)

4. The user can request that MetaFlux compute an initial set of biomass metabolites for the organism. MetaFlux will do so if the organism falls within the twelve taxonomic groups for which MetaFlux has defined biomass compositions, obtained from the experimental literature.

5. The user supplies an objective function and constraints on metabolite uptake and secretion. These can be based on experimental observations of the organism under study, or can be set to arbitrary values to explore the theoretical behavior of metabolism.

6. The MetaFlux gap filler identifies missing reactions, nutrients, and secretions that enable a model to be solved. It can be run on one compartment at a time to simplify the gap-filling process for eukaryotic organisms. MetaFlux tools for computing blocked reactions and dead-end metabolites identify potential errors and omissions in the metabolic network definition.

7. The reaction network, objective function, and constraints of the metabolic model are adjusted by the user until its predictions match experimental results.

When developing a metabolic model with MetaFlux, the reactions and metabolites within the model are derived from (and stored in) the PGDB. MetaFlux automatically generates the system of linear equations for the model from the PGDB. Thus, to modify the reactions within a model, the user edits the PGDB; to inspect the reactions and metabolites within a model, the user can query the PGDB using the plethora of Pathway Tools query and visualization operations. The entire PGDB/model can be published on a website using Pathway Tools, where, all reactions and pathways that utilize a given metabolite are listed on the Pathway Tools metabolite page for that compound; all model metabolites within a given molecular weight range can be found using metabolite searches; all reactions within a given cellular location or using a given set of reactants and products can be found using reaction searches.

Furthermore, compared to other modeling environments, a metabolic model stored within a PGDB contains extensive additional enriching information. Chemical structures within a

PGDB enable reaction mass and charge balancing. Chemical structures and reaction atom mappings aid users in understanding the chemistry of reaction transformations. Pathways arrange reactions into biologically meaningful groupings. Couplings between reactions, enzymes, and genes enable reasoning about the roles of multi-subunit complexes, isozymes, and gene knockouts. Regulatory information supports inferences about metabolic regulation. Model testing and validation are facilitated by PGDB storage of growth media and growth experiment results, and of gene-knockout experiment results.

Taken together, the Pathway Tools modeling environment—with its extensive tools for inspecting metabolic-model content and its enriching information—renders MetaFlux models significantly easier to understand, learn from, validate through inspection, reuse, and extend than models produced with other metabolic-modeling software environments.

10.2 Description of an FBA Model

The description of a flux-balance analysis (FBA) model is provided to MetaFlux partly via a text file called an FBA input file and partly via a PGDB. Typically, the FBA file specifies many parameters, but we will describe only the most important ones. Some parameters are only relevant for specific modes of MetaFlux, so that we will present these parameters when describing that mode. All parameter names end with a colon, whereas keyword options for parameters start with a colon. In the following explanation, we omit the colon character to reduce clutter.

10.3 Solving Mode

Solving mode computes flux values for the reactions in the metabolic model given four inputs: a set of nutrient compounds, a set of secreted compounds, a set of biomass metabolites that are synthesized by the cell, and a set of metabolic reactions. The first three inputs are supplied by the FBA file. The set of metabolic reactions are provided by the PGDB, but may be altered by the FBA file. For example, the following file describes a (very simple) FBA model for *E. coli*: **[Dan: Shouldn't MAL, FUM, and WATER have compartment specifications?]**

```
pgdb: ecoli

reactions:
metab-all           # Include all metabolic reactions of PGDB ecoli.
FUMHYDR-RXN         # Example of including a reaction according to its
                    # PGDB unique identifier.
```

```
mal -> fum + water  # Example of including a reaction by reaction equation.

biomass:

CYS[CCO-CYTOSOL] 0.0054
GLN[CCO-CYTOSOL] 0.2987
GLT[CCO-CYTOSOL] 0.2987
GLY[CCO-CYTOSOL] 0.3431

nutrients:

GLC[CCO-PERI-BAC]              :upper-bound 10.0
OXYGEN-MOLECULE[CCO-PERI-BAC] :upper-bound 1.0

secretions:
CARBON-DIOXIDE[CCO-PERI-BAC]
WATER[CCO-PERI-BAC]
```

Note that this example is meant to show the syntax of the file describing a model, it is not meant to show a working model. The `pgdb` parameter specifies the PGDB to be used in this model. The set of model reactions is specified by using the `reactions` parameter. The keyword `metab-all` specifies all metabolic reactions from the PGDB, whereas `FUMHYDR-RXN` specifies one reaction using a unique identifier. A reaction equation can be provided to describe a reaction that is to be present in the model, but is not present in the PGDB.

The set of biomass metabolites are specified by parameter `biomass`; each metabolite is specified either by metabolite name or unique identifier, plus a compartment identifier in square brackets, and an optional coefficient. Nutrients and secretions are provided in the same manner, but no coefficients are allowed. Upper and lower bounds can be provided to constrain nutrient uptake rates or secretion production rates.

Specifying the metabolites of the biomass reaction using groups is also possible. The use of groups present a better structure of the biomass reaction. For example,

```
biomass:

:group val
Charged-VAL-tRNAs[CCO-CYTOSOL]   0.423162
VAL-tRNAs[CCO-CYTOSOL]          -0.423162
:end-group
```

```
:group thr
Charged-THR-tRNAs[CCO-CYTOSOL]    0.253687
THR-tRNAs[CCO-CYTOSOL]           -0.253687
:end-group
```

These two groups, one named `val`, the other named `thr`, specify a relationship between a charged and an uncharged tRNA, for valine and threonine. Essentially, the groups enable gathering the compounds where one or more compounds, with negative coefficients, are needed to produce the other(s) with positive coefficients. In other words, these groups describe partial sub-reactions of the biomass reaction because the metabolites with negative coefficients can be considered as reactants, and the metabolites with positive coefficients can be considered as products of a reaction, which is not necessarily mass balanced. The negative and positive values do not need to be the same absolute values because these sub-reactions are allowed not to be mass balanced. Metabolites can be repeated in different groups, but repetition of a metabolite is not allowed outside groups. Groups have another use with the `try-biomass` parameter in development mode (see subsection 10.4).

10.3.1 Modeling a Community of Organisms

Almost all organisms live in a larger ecosystem. MetaFlux can solve a model describing an organism community by combining individual FBA input files for each organism in the community, plus one COM input file describing the organism interactions. A typical COM description lists the FBA models to be combined, optionally with the abundance of each organism in the community. It lists the compartments shared by the organisms — the compartments in which the secretions of some organisms can be used as nutrients by other organisms. For example, the following COM file specifies two models, the exchange of metabolites in the periplasmic space, and three nutrients that are provided to the community.

```
community-name: ecoli-ile

fba-files:
  ecoli-strain-A.fba :abundance 2
  ecoli-strain-B.fba :abundance 3

exchange-compartments: [CCO-PERI-BAC]

override-fba-nutrients: no

community-nutrients:

GLC[CCO-PERI-BAC]              :upper-bound 20
OXYGEN-MOLECULE[CCO-PERI-BAC] :upper-bound 10.0
```

AMMONIUM[CCO-CYTOSOL]

The objective function defined by MetaFlux for the FBA model of a community of organisms is the sum of the growth of the organisms multiplied by their abundances. A solution file will be produced that describes the growth of the community, and the growth rate for each organism. We are currently developing an alternative community modeling approach inspired by COMETS [90], in which each organism has an independent objective function, and dynamic FBA is used to create a temporal simulation.

10.4 Development Mode

The development mode can be used to create an FBA model or to discover what is wrong with a model that does not grow when growth is expected. The main parameters used in development mode are the four parameters `try-biomass`, `try-reactions`, `try-secretions` and `try-nutrients`.

When MetaFlux is provided with a list of metabolites for the `try-biomass` in development mode, the software tries adding these metabolites to the biomass reaction. That is, in development mode, MetaFlux computes the maximum number of `try-biomass` metabolites that can be produced by the model. In other words, MetaFlux will output as a solution the largest subset of metabolites it can produce as biomass. In the early phases of developing a model, typically the entire biomass reaction is specified by `try-biomass` and no metabolites are specified for the `biomass` parameter.

The `try-biomass` parameter cannot specify any metabolite with a negative coefficient not included in a group. This is required because any such metabolite could be used as a "free" nutrient for the organism. On the other hand, inside a group, negative coefficients are allowed because MetaFlux tries to produce the entire group of metabolites, not any one of them independently. A group of metabolites is supposed to form a cohesive unit where any metabolite specified with a negative coefficient is used to produce some other metabolites of that group. Therefore, declaring all metabolites that have negative coefficients in some groups enables the entire biomass reaction to be used as a try-biomass set, indicating for all metabolites which ones can or cannot be produced.

The `try-reactions` parameter can be used to try to add candidate reactions to the model to increase the number of `try-biomass` metabolites produced by the model. Therefore, the `try-reactions` parameters is typically used when at least one biomass metabolite is not produced. The candidate reactions are selected from MetaCyc. A single keyword, `metacyc-metab-all`, instructs MetaFlux to try all the metabolic reactions of MetaCyc. Alternatively, a list of candidate reactions can be specified by their unique identifiers. MetaFlux tries to produce as many biomass metabolites as possible from the `try-biomass` section, by adding as few reactions as possible from the `try-reactions` set. This computation is

performed using optimization as a Mixed-Integer Linear Program (MILP), which can be computationally expensive to solve. The amount of time to compute an optimal solution varies depending on the number of candidate reactions. It can take a few seconds to several hours, or even days.

MetaFlux has also a fast development mode that can be used with the `try-reactions` parameter only. It may run much faster than the general development mode and may provide different solutions. It uses an heuristic that does not necessarily provide an optimal solution. But optimality is not always the best possible biological solution. Please consult [91] for a more complete description of the advantages and disadvantages of fast development mode.

10.5 Knockout Mode

Knockout mode is used to computationally evaluate the impact of removing genes or reactions from an FBA model. This mode is used to predict essential genes of an organism for a given growth environment, and can also be used to evaluate the accuracy of a model if experimental gene-knockout data are available. MetaFlux can compute single, double, or higher numbers of simultaneous knockouts.

When run, MetaFlux solves the model without any knockouts, and then solves for each gene to knockout by deactivating the reactions associated with that gene. Note that a given gene may deactivate one, none, or several reactions, because some genes may have isozymes, or catalyze several reactions. Requesting a summary only of the results of the knockout or requesting that, beside the summary, a complete solution file be produced for each gene knockout is also possible. See the subsection 10.6.2 for more details on the solution files produced.

10.6 Outputs Generated by MetaFlux

Whichever mode is used to execute MetaFlux, the following output files are generated: the solution file that contains a description of the active reactions and metabolites used and produced, the log file to describe issues that may exist in the model, and a data file for the fluxes that can be displayed with the Cellular Overview map of Pathway Tools.

In solving mode, the main output is the computed optimal assignments of reaction fluxes. For an organism-community model, solution, log, and data files are produced for each organism.

In development mode, the outputs are the set of biomass metabolites that can be produced; the try nutrients utilized (if any); the try secretions produced (if any); the reactions that actively carry flux; and a minimal list of suggested reactions to be added to the model, and reactions whose directions are reversed (if any) to produce otherwise unproducible biomass metabolites. Development mode also identifies which biomass metabolites could not be

produced by the model (if any), despite the additions from the try sets.

10.6.1 Reports Generated by MetaFlux

The set of reactions specified by the PGDB and the FBA input file may contain reactions that cannot, or that should not, be used in a model. MetaFlux checks each reaction to ensure that its inclusion in the model will not invalidate the model. In particular, all reactions are verified to be mass and charge balanced. Reactions that are unbalanced, or cannot be shown to be balanced, are not included in the model, but are listed in the log file.

Another step in model execution is to instantiate the generic reactions of a model according to the compounds available in the PGDB. A generic reaction has some compound classes as products or reactants (e.g., "sugars"). Each computed reaction instantiation is added to the model. If a generic reaction instantiation was not possible (for example, when which compound instances should be used for classes on each side of a reaction it is ambiguous) then such generic reactions are listed in the log file.

Blocked reactions. MetaFlux analyzes the network of reactions specified by the PGDB plus the FBA input file to detect if some reactions are blocked. A blocked reaction is a reaction that can never have a positive flux, given the nutrients, secretions, biomass metabolites, and reactions specified for a particular model execution. That is, blocked reactions are a function not only of the network, but of the cellular growth conditions. A blocked reaction R has at least one reactant M that is not produced by any reaction in the model, and that is not provided as a nutrient; or R has a product M that is not used as a reactant by any other reaction and that is not secreted, and is not specified as a biomass metabolite with a positive coefficient. These reactions are the *basic blocked reactions*, and the metabolites that caused basic blocked reactions to be blocked are called *basic blocking metabolites*.

Blocked reactions can never carry flux because in steady-state modeling, the producing and consuming fluxes for every metabolite must be balanced. But the preceding metabolites M could not have balanced fluxes if a blocked reaction R carried flux, because according to the preceding definition, M must have either zero reactions that produce it or zero reactions that consume it.

Additional blocked reactions can be identified by eliminating basic blocked reactions from the model, causing more reactions to become basic blocked reactions. This process of removing basic blocked reactions from the model is repeated until no more reactions become blocked. The detection of blocked reactions is done before the linear solver is called (that is, this detection does not depend on the fluxes of reactions, but is a static evaluation of the model).

The set of blocked reactions is listed in the log file, grouped by basic blocking metabolite (the reactions are also grouped by pathways). The basic blocking metabolites are, in a sense, the root causes of blocked reactions, thus their identification is quite valuable for model debugging.

10.6.2 The MetaFlux Solution File

The content of the solution file generated by MetaFlux depends on the mode of MetaFlux. In solving mode, the solution file contains the uptake rates of the nutrients, the production rates of the biomass metabolites and secretions, and the fluxes of the model reactions. If no growth is obtained (that is, the flux of the biomass reaction is zero, or very near zero), this condition is stated in the solution file.

When solving a community model, a community solution file is generated that contains a summary of the result, and a complete solution file is produced for each organism in the community. The summary lists many values: the growth rate of each model, the compartments that were involved in the exchange of metabolites between the organisms; the secretions used as nutrients by some organisms with their flux; the secretions that were not used by any organisms; and the secretions that were not produced by any organism of the community. In essence, the summary gives the overall view of the exchange of metabolites between the organisms and the rate of growth of each organism.

In development mode, the solution file contains essentially what is given in solving mode, but the flux values are only meaningful as to whether they are zero or non-zero. For example, if a `try-biomass` parameter was specified with some metabolites, the metabolites that could or could not be produced are identified. This is different from the solving mode, where only one biomass metabolite that cannot be produced will result in no growth, but with no indication of which metabolites could not be produced.

In knockout mode, a solution file with a summary of the results is produced. The summary includes the growth of the model with no knockouts, followed by a list of every knockout that was requested, based on genes and/or reactions to knockout, and the resulting flux for the biomass reaction. Requesting a complete solution file for each knockout performed is also possible.

10.6.3 Painting MetaFlux Fluxes on the Cellular Overview

MetaFlux also generates a data file that can be used with the Cellular Overview of Pathway Tools (see Section 5.4). The graphical user interface of MetaFlux enables the user to click one button to invoke the Cellular Overview with the data from that file. The Cellular Overview displays all the reactions, grouped by pathways, of the organisms, and highlights, with colors indicative of flux values, the reactions that have positive fluxes. This overview enables the user to visually assess the metabolic activities of a model solution.

For a community of organisms, a data file is generated for each individual organism, and the graphical user interface gives direct access to each individual Cellular Overview of the organisms involved in the community.

11 Survey of Pathway Tools Compatible Databases

According to user surveys, Pathway Tools users have created 9,800 PGDBs for organisms from all domains of life, in addition to the 5,500 PGDBs available from BioCyc.org. This section summarizes those user-created PGDBs that are publicly available.

With highly curated PGDBs available for many important organisms, it is not clear why users would consider using the uncurated (and therefore lower quality) pathway DBs available for these same organisms from other pathway DB providers. For example, consider the highly curated AraCyc pathway database for *Arabidopsis thaliana* [92, 93]. AraCyc contains mini-review summaries for enzymes and metabolic pathways; thousands of literature references; evidence codes for enzyme functions and metabolic pathways (indicating which pathways are supported by experimental evidence); and information on enzyme subunit structure, activators, inhibitors, and cofactors. In contrast, KEGG data on *Arabidopsis* contains none of the preceding information. In addition, AraCyc curators have carefully refined the metabolic reactions and pathways present in AraCyc, such as to remove false-positive computational predictions, and to add *Arabidopsis* reactions and pathways from the biomedical literature to AraCyc. Although KEGG updates its reference pathway map diagrams periodically to contain new pathways and reactions from different organisms, the KEGG approach of computationally coloring reactions within pathway maps based on the presence of enzymes for those reactions within a genome results in significant ambiguity. If AraCyc curators are reasonably certain that a reaction or pathway is absent from *Arabidopsis*, they remove it from the database. The KEGG model does not allow such removal, so whether an uncolored reaction is truly absent from an organism or whether the gene for its enzyme has not yet been identified in the genome is never clear within KEGG. This situation results in a real conundrum for a scientist who wishes to assemble the list of reactions likely to be present in *Arabidopsis* from KEGG, because is no way exists to distinguish the many uncolored reactions that are likely present but for which no gene has been identified, from the many uncolored reactions that are clearly known to be absent from *Arabidopsis* (which curators have deleted from AraCyc).

Available PGDBs include those in Tables 1–4, with curated PGDBs in bold.

Hosting Site	**Organism**	**Citations, URLs**
BioCyc	***Bos taurus* (CattleCyc)**	[94, 95]
FlyCyc	***Drosophila melanogaster* (FlyCyc)**	[96]
BioCyc	***Homo sapiens* (HumanCyc)**	[97, 98]
MouseCyc	***Mus musculus* (MouseCyc)**	[99]

Table 1: **PGDBs for animals.**

To facilitate sharing of PGDBs among multiple users, we have created a PGDB Registry that enables peer-to-peer sharing. PGDB sharing is desirable because a user whose own

Hosting Site	Organism	Citations, URLs
CoffeaCyc	*Coffea canephora*	[100]
MediCyc	***Medicago truncatula***	[101, 102]
Gramene	*Arabidopsis thaliana*	[103]
Gramene	*Brachypodium distachyon*	[104]
Gramene	*Coffea canephora*	[105]
Gramene	*Medicago truncatula*	[106]
Gramene	*Oryza sativa*	[107]
Gramene	*Populus trichocarpa*	[108]
Gramene	*Sorghum bicolor*	[109]
Gramene	*Solanum lycopersicum*	[110]
Gramene	*Solanum tuberosum*	[111]
Gramene	*Zea mays*	[112]
sol genomics network [113]	*Capsicum annuum*	[114]
sol genomics network	*Nicotiana benthamiana*	[115]
sol genomics network	*Nicotiana tabacum*	[116]
sol genomics network	*Petunia hybrida*	[117]
sol genomics network	*Solanum lycopersicum*	[118]
sol genomics network	*Solanum tuberosum*	[119]
Plant Metabolic Network [120]	**PlantCyc**	[120]
Plant Metabolic Network	***Arabidopsis thaliana*** **(AraCyc)**	[92, 93]
Plant Metabolic Network	*Brachypodium distachyon*	[121]
Plant Metabolic Network	*Brassica rapa*	[122]
Plant Metabolic Network	*Carica papaya*	[123]
Plant Metabolic Network	*Chlamydomonas reinhardtii*	[124]
Plant Metabolic Network	*Manihot esculenta*	[125]
Plant Metabolic Network	*Hordeum vulgare*	[126]
Plant Metabolic Network	*Glycine max*	[127]
Plant Metabolic Network	*Panicum virgatum*	[128]
Plant Metabolic Network	*Physcomitrella patens*	[129]
Plant Metabolic Network	*Populus trichocarpa*	[130]
Plant Metabolic Network	*Oryza sativa*	[131]
Plant Metabolic Network	*Selaginella moellendorffii*	[132]
Plant Metabolic Network	*Setaria italica*	[133]
Plant Metabolic Network	*Sorghum bicolor*	[134]
Plant Metabolic Network	*Vitis vinifera*	[135]
Plant Metabolic Network	*Zea mays*	[136]

Table 2: **PGDBs for plants.**

Hosting Site	Organism	Citations, URLs
Ano2Cyc	***Anopheles gambiae***	[137]
CandidaCyc	***Candida albicans***	[138]
DictyCyc	***Dictyostelium discoideum***	[139]
LeishCyc	***Leishmania major***	[140]
PchrCyc	***Penicillium chrysogenum***	[141]
YeastCyc	***Saccharomyces cerevisiae*** **(YeastCyc)**	[142]
ThapsCyc	***Thalassiosira pseudonana***	[143]
TrypanoCyc	***Trypanosoma brucei***	[144]
BioCyc	**8 PGDBs for *Plasmodium, Cryptosporidium,* and *Toxoplasma***	[95]

Table 3: **PGDBs for other eukaryotes.**

Hosting Site	Organism	Citations, URLs
BioCyc	***Escherichia coli*** **(EcoCyc)**	[8, 145]
BioCyc	***Bacillus subtilis***	[146]
BioCyc	***Candidatus cardinium***	[147]
BioCyc	***Candidatus portiera aleyrodidarum***	[148]
BioCyc	***Candidatus Evansia muelleri***	[149]
BioCyc	***Clostridium saccharoperbutylacetonicum***	[150]
BioCyc	***Listeria monocytogenes***	[151]
BioCyc	***Methylosinus trichosporium***	[152]
BioCyc	***Peptoclostridium difficile***	[153]
BioCyc	***Streptomyces coelicolor***	[154]
PseudoCyc	***Pseudomonas aeruginosa*** **(PseudoCyc)**	[155]
BioCyc	5500 PGDBs	[9, 95]
MicroCyc	1702 PGDBs	[156]
Taxonomically Broad EST Database	49 PGDBs	[157, 158]
BeoCyc	23 PGDBs	[159]

Table 4: **PGDBs for microbes.**

computer has a copy of a PGDB can use Pathway Tools functionality that would not be available through a remote Pathway Tools web server, such as functionality that exists in desktop mode only, and such as comparative operations. Comparative analysis of two or more PGDBs is possible only when they are loaded into the same instance of Pathway Tools.

The PGDB Registry uses a server maintained by SRI that tracks the locations of available PGDBs that PGDB authors have registered for downloading. The author of a PGDB can register that PGDB by using a command within Pathway Tools that creates an entry for the PGDB in the Registry server, and places the PGDB on an FTP or HTTP server of

the author's choosing. Users who want to download a PGDB from the Registry can view available PGDBs by using a web browser (see URL [160]) or using Pathway Tools itself. With a few mouse clicks, a user can download a PGDB from the registry using Pathway Tools.

12 Comparison of Pathway Tools with Related Software Environments

Pathway Tools stands out with respect to related software tools in the breadth of the functionality and the high level of integration that it provides. It addresses a very large number of use cases. And it provides schema, visualization, and editing support for an unusually large number of data types in addition to metabolic and signaling pathways, including chromosomes, genes, enzymes, transporters, regulatory networks, and compounds. The following comparison is organized roughly according to the use cases presented in Section 2.

12.1 Development, Visualization, and Web Publishing of Organism Specific Databases

12.1.1 Metabolic Pathway Information

Other software systems for managing metabolic pathway information are KEGG [161, 162]; The SEED [163]; GenMAPP [164, 165]; PathCase [166, 167]; VisANT [168]; and Reactome [169]. KEGG, Reactome, and GenMAPP employ static, pre-drawn pathway diagrams, a model that does not scale to produce custom pathway diagrams for thousands of *different* pathways in different organisms. Nor can the static approach produce multiple views of a given pathway at different levels of detail, as can the Customize Pathway option in Pathway Tools that enables the user to choose exactly which graphical elements (e.g., gene names, EC numbers, metabolite structures, activators and inhibitors) appear in the pathway diagram, and to color a pathway diagram with omics data.

PathCase and VisANT do have pathway layout capabilities, but the resulting diagrams bear little resemblance to those found in the biomedical literature, nor are they particularly compelling visually. Their diagrams can be customized in various ways (e.g., choice of layout algorithm, data overlays, and selective compaction of nodes), but they do not offer the same types of customization or multiple-detail views offered by Pathway Tools.

Cytoscape [170] is a general tool for displaying biological networks that embodies the philosophy that general graph layout techniques can satisfactorily depict any biological network. Although the Cytoscape layout algorithms are a terrific fit for display of protein interaction maps, we assert that the general layout algorithms do not produce useful results for

metabolic pathways. We believe that superior visualization results are obtained when layout algorithms are specifically tailored to metabolic pathways. For example, Pathway Tools provides separate layout algorithms for circular, linear, and tree-structured pathways to make the structure of those pathways stand out prominently to the biologist. Biologists developed their pre-computer depictions of metabolic pathways for important reasons, namely to clearly present the topology of a pathway.

SEED models lack a unifying pathway visualization framework—they use KEGG pathway diagrams, often supplemented with simple text-based diagrams and/or images imported from other sources, with limited interactivity beyond hyperlinks.

GenMAPP, Reactome, VisANT, KGML-ED [171], and PathCase have pathway-editing capabilities. KEGG and SEED lack pathway editing, meaning that users cannot introduce new organism-specific pathways, nor modify a reference pathway to customize it to a specific organism, thus eliminating the possibility of removing erroneous reaction steps from a pathway, or of adding missing reactions to a pathway.

KEGG Atlas [172] has a comparable analog of our Cellular Overview diagram (which we introduced in 1999 [173, 47]). The KEGG Atlas diagram was constructed by hand and provides a single set of overview metabolic maps for all organisms in KEGG, as opposed to the organism-specific overviews that Pathway Tools generates through advanced layout algorithms. The KEGG diagram is not as flexibly queryable or interactive as the Pathway Tools diagram, and it also lacks semantic zooming for adding increasing details. A notable improvement occurred with Pathway Projector [174]. It relies directly on KEGG for the overall overview diagram and other data, but graphically, it is enhanced, adding EC numbers and genes to reaction edges.

Reactome includes a different type of high-level overview diagram called Fireworks, which shows a hierarchical organization of nodes representing metabolic and signaling pathways, mostly the latter. Users can navigate to pathways of interest with this diagram, and highlight pathways that contain objects such as compounds or genes that are supplied in an analysis interface. Fireworks can zoom in until pathway names are shown. From there, one can click to a specific pathway display for the details. However, Fireworks does not depict individual reactions or metabolites; therefore, it is not a metabolic map diagram.

Tools such as VisANT and Cytoscape can be used to display very large networks, including the entire metabolic network for an organism, and these diagrams can be used to visualize omics data and interacted with in other ways, but given that such diagrams are laid out using general-purpose network layout algorithms, individual pathways are not likely to be recognizable.

No other general tool can show reactions with atom-mapping information.

Computationally generated diagrams for electron transfer reactions and pathways, as well as for transport reactions, with their depiction of compartments at a cellular membrane, are

unique features of Pathway Tools.

12.1.2 Genome and Proteome Information

Most model-organism DB tools include genome browsers and gene pages. A representative sample of such systems includes GBrowse [175, 176]; JBrowse [177]; IMG [178]; Entrez Genome [179]; the UC Santa Cruz (UCSC) Genome Browser [180]; Ensembl [181]; and PATRIC [182].

The Pathway Tools genome browser uses a different approach than the others. First, all the other genome browsers depict the chromosome as a single horizontal line, with tracks information below it. Pathway Tools uses this approach when the user explicitly enables tracks, but when tracks are not enabled, a multi-line wrapped display is used to present the most possible information in the available display window. Second, all the other genome browsers use tracks to depict coding regions, promoters, and other genome features. Although Pathway Tools does provide a tracks capability, its assumption is that consensus curated information on coding region extents, promoters, terminators, and transcription-factor binding sites should be depicted "inline" on the multi-line chromosome. This approach results in a much more information-dense display.

GBrowse is highly customizable through the addition of data tracks. Wrapped multiline displays that will make full use of screen real estate are not available. In GBrowse, the semantics of semantic zooming has to be configured but provides flexibility. GBrowse can provide comparative genomics with additional tools like SynView [183] or GBrowse_syn [184].

JBrowse is a JavaScript-based genome browser that could become a replacement for GBrowse. Its main innovation is to rely on the web browser to perform most of the processing and drawing, making it fast and smooth. PATRIC uses JBrowse.

The IMG genome browser (or "chromosome viewer", in their parlance) is very basic, and does not support zooming. Some views are shown as wrapped multiline displays, and several gene coloring schemes can be selected, but no data tracks are offered. For comparative genomics, three different, unrelated synteny viewers are available.

The Entrez Genome sequence viewer has no real semantic viewing other than showing the DNA sequence at maximum zoom level. It offers support for adding extra tracks and an alignment view for comparative genomics.

The UCSC Genome Browser is mainly geared towards eukaryotes and has an extensive support for addition of data tracks. But no real semantic viewing is available other than showing the DNA sequence at maximum zoom level.

Ensembl is focused on eukaryotes, similarly to UCSC, and also has extensive support for extra tracks. It offers no real semantic viewing, other than showing the DNA sequence at maximum zoom level.

Most sequenced genomes are circular, but Pathway Tools seems to have the only genome browser that can seamlessly depict the junction between the first and last base-pair to show a contiguous view. All other genome browsers seem to artificially linearize all chromosomes. Also, Pathway Tools seems to be unique in providing a compact overview of the entire genome (the "genome overview"), which enables coloring each gene with omics data.

12.1.3 Regulatory Networks

A number of bioinformatics databases include regulatory network information; however, the majority of these databases and their associated software environments can represent information on transcription-factor-based regulation only, such as RegTransBase [185]; TRANSFAC [186]; CoryneRegNet [187]; ProdoNet [188]; and DBTBS [189]. The exception is RegulonDB [190], which can also capture RNA-based regulation, including riboswitches, attenuators, and small-RNA regulators.

We are not aware of tools comparable to the Regulatory Overview in being able to display and interrogate large, complete cellular regulatory networks, although CoryneRegNet and ProdoNet display smaller regulatory networks. CoryneRegNet also displays omics data onto its regulatory network diagrams.

12.1.4 Query Tools

Other organism DBs provide a subset of the three tiers of queries provided by Pathway Tools (quick search, object-specific searches, and Structured Advanced Query Page). Virtually all provide a quick search. Sites providing particularly extensive object-specific searches are FlyBase [191]; Mouse Genome Informatics [192]; EuPathDB [193]; and BioMart [194]. BioMart is used by bioinformatics DBs, including WormBase, Rat Genome Database, UniProt, Reactome, and Galaxy. Its underlying query language is Perl, using the BioMart libraries. However, none of the preceding systems provides the query power of the Pathway Tools SAQP. For example, BioMart does not allow the user to construct arbitrary queries that perform joins (queries that combine multiple data types); it provides only the "and" logical operator (the "or" operator is not available); and it includes only a limited form of "not." Biozon [195] (`biozon.org`) provides a web interface that provides fairly complex querying, including join operations, over several biological DBs.

12.2 Extend Genome Annotations with Additional Computational Inferences

KEGG, Model SEED, and Reactome are the only other tools that can predict pathways from genome data. The pathway hole filler and transport inference parser tools are unique

to Pathway Tools. Many genome-annotation pipelines include operon predictors.

12.3 Analysis of Omics Data

Kono et al. introduced a tool based on the Google Maps API for painting omics data onto an enhanced KEGG Atlas map, called Pathway Projector [174]. This tool does not produce animations as our omics viewers do. However, it can depict time-series expression data as small histograms that are reminiscent of Pathway Tools' omics pop-ups. It also provides more powerful highlighting options than KEGG Atlas. Sequence-based search and highlighting is available via KEGG BLAST, although it can take minutes. Users can manually annotate the map with custom markers and lines, which can be exported via XML and shared with others.

KEGG Atlas [172] enables some limited highlighting, but mostly by using KEGG identifiers. No animation or omics pop-ups seem to be available.

Reactome [169] can paint omics data onto the Fireworks hierarchical pathway overview. Zooming to high detail reveals the pathway names, but zooming does not go to the reaction level. Showing time-series data as animations is possible. No omics pop-ups seem to be available.

GenMapp [196]; VitaPad [197]; VisANT; and ArrayXPath [198] paint omics data onto single pathways, rather than onto a full metabolic overview.

12.4 Metabolic Modeling

Other well-known tools providing constraint-based metabolic modeling include COBRA [88] and Model SEED [199]. Many other FBA tools exist, but a detailed comparison with all of them is beyond the scope of this paper. Lakshmanan et al. provide a comparison among 19 FBA tools, including MetaFlux (see their Tables 2, 3, and 4) [200]. This paper was first published in 2012 and therefore covers an old version of MetaFlux. Additions to MetaFlux since that time include expansion to the Windows platform, import of SBML files, solving of models for organism communities, and support for dynamic FBA. Other tools have probably developed new capabilities as well. We now examine some of the issues raised by Lakshmanan et al. in more detail.

Our design of MetaFlux has emphasized ease of use and acceleration of the very slow model development times of 12–18 months stated in some FBA publications. At SRI we have reliably been able to create FBA models for bacteria and fungi in 3–4 weeks (albeit to a lower level of validation than some published models). We note that in another recent survey of metabolic-modeling software [201], Pathway Tools was the only software package recommended for users new to modeling.

Elements of the MetaFlux approach that speed model development are as follows. (a) Use of MetaFlux does not require programming ability as some other tools do. (b) Lakshmanan mentions the collection of metabolic reactions for an organism at the start of all metabolic modeling projects in passing, without noting the strong link between model accuracy and the completeness of the initial reaction set. The PathoLogic enzyme name matcher (see Section 3.2) combined with the extensive reaction database of MetaCyc [9] provide a powerful resource for reactome prediction. Furthermore, when metabolic pathway prediction follows reactome prediction, it produces a more complete metabolic network, because it imports all reactions in a pathway, even those that lack enzyme assignments, thus simplifying the later gap-filling step. Furthermore, the fact that Model SEED contains approximately 1/3 as many metabolic pathways as does MetaCyc raises a major question in the use of Model SEED for pathway reconstruction.[4] (c) MetaFlux has a more powerful gap filler than any other tool: it can gap fill not only reactions but also nutrients and secretions, and it can identify which biomass metabolites cannot be synthesized by the model, a critical step in helping the user discover which aspect of the network is incomplete. This is a feature not found in other tools to our knowledge. (d) MetaFlux computes blocked reactions and metabolites, reporting to the user the basic blocking metabolites that are the root causes of model blockages. (e) MetaFlux computes the reaction balance from the reaction equations plus chemical structure data to ensure that unbalanced reactions are prohibited from inclusion in a model.

Many other tools use the GLPK solver (some not exclusively). MetaFlux uses the SCIP solver [83], which has a much faster Mixed-Integer Linear Programming (MILP) solver, needed for gap filling, and is free for research use by academics. The SCIP MILP solver is fast, although not as fast as commercial solvers.

The comment by Lakshmanan et al. on p8 that "none of the [network visualization] tools can handle large-scale models" clearly does not apply to MetaFlux, as its cellular overview diagram does handle genome-scale models.

Lakshmanan et al. suggest (p10) that more tools should consider linking with a biological model storage DB. Pathway Tools has taken this approach to an extreme, as described in Section 10.1.

13 Limitations and Future Work

Here, we summarize limitations of Pathway Tools, organized by use case. Some of these limitations are being addressed in current research; many of the others will be addressed in future work.

[4]This number was obtained by manually reviewing the 1,320 subsystems present in the February 2015 version of Model SEED, then removing all subsystems not related to metabolism, then removing subsystems for individual metabolic enzymes. This process yielded 730 subsystems that correspond to metabolic pathways, compared to the 2,300 pathways in the February 2015 version of MetaCyc.

Development of Organism Specific Databases. Pathway Tools has an emphasis on prokaryotic biology, although over time we have added, and plan to add, more support for eukaryotic biology. One remaining limitation is that although the software can capture many types of prokaryotic regulation, we have not attempted comprehensive coverage of eukaryotic regulation. Another limitation is that the editing tools within Pathway Tools are not web based, but require installation of Pathway Tools on every computer that will be used for editing.

Visualization and Web Publishing of Organism Specific Databases. In recent years we have made significant progress in making the capabilities of the web and desktop modes of the Pathway Tools Navigator as similar as possible. For example, many cellular overview capabilities previously present only in desktop mode are now also available in web mode. However, not all capabilities of Pathway Tools are available in both the web and desktop modes. For example, many comparative tools function in web mode only, whereas all aspects of PathoLogic and the editing tools are available in desktop mode only.

Analysis of Omics Data. Pathway Tools is not a general-purpose environment for analysis of omics data. Our assumption is that scientists will use one of the many other software packages for the early stages of omics data analysis (such as normalization), and will provide the output of those analyses to Pathway Tools for display with the omics viewers.

Analysis of Biological Networks. We would like to see many additional network analysis tools present within Pathway Tools, such for computing the scaling properties of metabolic networks [202], and functional modules within metabolic networks [203].

Comparative Analysis of Organism-Specific Databases and Metabolic Engineering. We are not aware of striking limitations in these areas.

14 Summary

Pathway Tools treats a genome as far more than a sequence and a set of annotations. Instead, it links the molecular parts list of the cell both to the genome and to a carefully constructed web of functional interactions. The Pathway Tools ontology defines an extensive set of object attributes and object relations that enables representing a rich conceptualization of biology within a PGDB, along with enabling querying and manipulation by the user. Furthermore, a PGDB can be transformed into a quantitative metabolic model for the organism.

Pathway Tools provides a broad range of functionality. It can manipulate genome data, metabolic networks, and regulatory networks. For each datatype, it provides query, visualization, editing, and analysis functions. It provides model-organism database development capabilities, including computational inferences that support fast generation of comprehensive databases, editors that enable refinement of a PGDB, web publishing, and comparative analysis. A family of curated PGDBs has been developed using these tools for important

model organisms.

The software also provides visual tools for analysis of omics datasets, and tools for the analysis of biological networks.

15 Software Availability

Pathway Tools runs on Macintosh, Windows, and Linux. It is freely available to academic and government researchers; a license fee applies to commercial use. See `http://BioCyc.org/download.shtml`.

Funding

This work was supported by grants GM75742, GM080746, and GM077678 from the National Institutes of Health. The contents of this article are solely the responsibility of the authors and do not necessarily represent the official views of the National Institutes of Health.

Acknowledgments

Pathway Tools has benefited from advice, input, and contributions from many scientists during its lifetime. We particularly wish to recognize contributions from Ian Paulsen, Robert Gunsalus, Monica Riley, John Ingraham, Jean-Francois Tomb, and Peifen Zhang. Lukas Mueller developed PerlCyc and has provided many helpful suggestions. Thomas Yan developed JavaCyc. Tomer Altman developed RCyc. Jeremy Zucker developed the SBML generation module, and contributed many other ideas. Christos Ouzounis was a co-developer of the original metabolic pathway prediction algorithm, contributed an early version of the import/export system, and has been a source of much sound advice.

Key Points

- The Pathway Tools software is a comprehensive environment for creating model organism DBs that span genome information, metabolic pathways, and regulatory networks.
- Pathway Tools inference capabilities include prediction of metabolic pathways, prediction of metabolic pathway hole fillers, inference of transport reactions from transporter functions, and prediction of operons.

- Its metabolic modeling capabilities include flux-balance analysis modeling for individual organisms and organism communities, with model gap filling and the ability to model gene knockouts.
- Pathway Tools provides interactive editing tools for use by database curators.
- Omics data analysis tools paint genome-scale datasets onto a complete genome diagram, complete metabolic network diagram, and complete regulatory network diagram.
- Other tools include comparative analysis operations, dead-end metabolite and blocked-reaction analysis of metabolic networks, and metabolic route searching.

References

[1] Karp P, Paley S, Krummenacker M et al. Pathway Tools version 13.0: Integrated Software for Pathway/Genome Informatics and Systems Biology. Brief Bioinform 11:40–79, 2010. URL `http://bib.oxfordjournals.org/cgi/content/abstract/bbp043`, doi: 10.1093/bib/bbp043.

[2] Paley S and Karp P. The Pathway Tools Cellular Overview Diagram and Omics Viewer. Nuc Acids Res 34:3771–8, 2006. URL `http://nar.oxfordjournals.org/cgi/content/full/34/13/3771`.

[3] Karp P, Paley S, and Romero P. The Pathway Tools Software. Bioinformatics 18:S225–S232, 2002.

[4] Karp P and Paley S. Integrated Access to Metabolic and Genomic Data. Journal of Computational Biology 3:191–212, 1996. URL `http://www.ai.sri.com/pubs/papers/Karp96:Integrated/document.ps`.

[5] Karp P, Latendresse M, Paley S et al. Pathway Tools version 19.0 update: Software for Pathway/Genome Informatics and Systems Biology. Brief Bioinform 2015. URL `http://bib.oxfordjournals.org/content/early/2015/09/25/bib.bbv079.abstract?keytype=ref&ijkey=IUo8YbQG2gbykM6`, doi: 10.1093/bib/bbv079.

[6] Ouzounis C and Karp P. Global Properties of the Metabolic Map of *Escherichia coli*. Genome Research 10:568, 2000.

[7] Karp P. Pathway Databases: A Case Study in Computational Symbolic Theories. Science 293:2040–4, 2001.

[8] Keseler I M, Mackie A, Peralta-Gil M et al. EcoCyc: Fusing model organism databases with systems biology. Nucleic Acids Res 41(Database issue):D605–12, 2013.

[9] Caspi R, Altman T, Billington R et al. The MetaCyc database of metabolic pathways and enzymes and the BioCyc collection of Pathway/Genome Databases. Nuc Acids Res 42:D459–71, 2014. URL http://nar.oxfordjournals.org/content/42/D1/D459.

[10] Costanzo M C, Engel S R, Wong E D et al. Saccharomyces genome database provides new regulation data. Nucleic Acids Res 42(Database issue):D717–25, 2014.

[11] Aurrecoechea C, Heiges M, Wang H et al. ApiDB: Integrated Resources for the Apicomplexan Bioinformatics Resource Center. Nuc Acids Res 35:D427–D430, 2007.

[12] Harris T W, Baran J, Bieri T et al. WormBase 2014: new views of curated biology. Nucleic Acids Res 42(Database issue):D789–93, 2014.

[13] Eppig J T, Blake J A, Bult C J et al. The Mouse Genome Database (MGD): facilitating mouse as a model for human biology and disease. Nucleic Acids Res 43(Database issue):D726–36, 2015.

[14] Krummenacker M, Paley S, Mueller L et al. Querying and computing with BioCyc databases. Bioinformatics 21:3454–5, 2005. URL http://bioinformatics.oxfordjournals.org/cgi/reprint/21/16/3454.

[15] Lee T, Pouliot Y, Wagner V et al. BioWarehouse: A bioinformatics database warehouse toolkit. BMC Bioinformatics 7:170, 2006. URL http://www.biomedcentral.com/1471-2105/7/170.

[16] Yeh I, Hanekamp T, Tsoka S et al. Computational analysis of *Plasmodium falciparum* metabolism: Organizing genomic information to facilitate drug discovery. Genome Research 14(5):917–24, 2004.

[17] Stephanopoulos G, Aristidou A, and Nielsen J. Metabolic Engineering: Principles and Methodologies. Academic Press, San Diego, CA, 1998.

[18] Karp P, Chaudhri V K, and Paley S M. A collaborative environment for authoring large knowledge bases. J Intelligent Information Systems 13:155–94, 1999. URL http://www.ai.sri.com/pkarp/pubs/99jiis.pdf.

[19] Paley S and Karp P. Evaluation of Computational Metabolic-Pathway Predictions for *H. pylori*. Bioinformatics 18(5):715–24, 2002.

[20] Dale J M, Popescu L, and Karp P D. Machine Learning Methods for Metabolic Pathway Prediction. BMC Bioinformatics 11:15, 2010.

[21] Karp P D, Latendresse M, and Caspi R. The Pathway Tools pathway prediction algorithm. Stand Genomic Sci 5(3):424–429, 2011.

[22] Romero P and Karp P. Using Functional and Organizational Information to Improve Genome-Wide Computational Prediction of Transcription Units on Pathway/Genome Databases. Bioinformatics 20:709–17, 2004.

[23] Green M and Karp P. A Bayesian method for identifying missing enzymes in predicted metabolic pathway databases. BMC Bioinformatics 5(1):76, 2004. URL `http://www.biomedcentral.com/1471-2105/5/76`.

[24] Lee T, Paulsen I, and Karp P. Annotation-based inference of transporter function. Bioinformatics 24:i259–67, 2008. URL `http://bioinformatics.oxfordjournals.org/cgi/content/full/24/13/i259`.

[25] Latendresse M, Malerich J, Travers M et al. Accurate Atom-Mapping Computation for Biochemical Reactions. J Chem Inf Model 2012.

[26] Jankowski M D, Henry C S, Broadbelt L J et al. Group contribution method for thermodynamic analysis of complex metabolic networks. Biophys J 95(3):1487–99, 2008.

[27] Alberty R A. Thermodynamics of Biochemical Reactions. Wiley InterScience, 2003.

[28] Funahashi A, Morohashi M, and Kitano H. CellDesigner: a process diagram editor for gene-regulatory and biochemical networks. Biosilico 1(5):159–62, 2003.

[29] Marvin Chemical Editor. `http://www.chemaxon.com/products/marvin/`.

[30] Ceroni A, Dell A, and Haslam S M. The GlycanBuilder: a fast, intuitive and flexible software tool for building and displaying glycan structures. Source Code Biol Med 2:3, 2007.

[31] Damerell D, Ceroni A, Maass K et al. The GlycanBuilder and GlycoWorkbench glycoinformatics tools: updates and new developments. Biol Chem 393(11):1357–1362, 2012.

[32] Green M and Karp P. The Outcomes of Pathway Database Computations Depend on Pathway Ontology. Nuc Acids Res 34:3687–97, 2006. URL `http://nar.oxfordjournals.org/cgi/content/full/34/13/3687`.

[33] Bochner B R, Gadzinski P, and Panomitros E. Phenotype Microarrays for High-Throughput Phenotypic Testing and Assay of Gene Function. Genome Research 11(7):1246–1255, 2001.

[34] Bochner B. Global phenotypic characterization of bacteria. FEMS Microbiol Rev 33(1):191–205, 2009.

[35] et al D F. The minimum information about a genome sequence (MIGS) specification. Nature Biotechnology 26:541–547, 2008.

[36] Karp P, Paley S, Krieger C et al. An Evidence Ontology for use in Pathway/Genome Databases. In R Altman and T Klein, editors, Proc Pacific Symposium on Biocomputing, 190–201. World Scientific, Singapore, 2004. URL http://www.ai.sri.com/pkarp/pubs/04psb-evidence.pdf.

[37] Pathway Tools Evidence Ontology. http://bioinformatics.ai.sri.com/evidence-ontology/.

[38] Pathway Tools Cell Component Ontology. http://bioinformatics.ai.sri.com/CCO/.

[39] Consortium G O. The Gene Ontology project in 2008. Nuc Acids Res 36:D440–4, 2008.

[40] Bairoch A, Apweiler R, and Wu C. UniProt Knowledgebase User Manual, 2009. URL http://www.expasy.ch/sprot/userman.html.

[41] Comparison of BioCyc Desktop Mode and Web Mode. http://biocyc.org/desktop-vs-web-mode.shtml.

[42] Alashqur A M, Su S Y W, and Lam H. OQL: A query language for manipulating object-oriented databases. In VLDB '89: Proceedings of the 15th International Conference on Very Large Data Bases, 433–442. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1989.

[43] Pathway Tools Data-File Formats. http://bioinformatics.ai.sri.com/ptools/flatfile-format.html.

[44] SOLR Website. http://lucene.apache.org/solr/resources.html.

[45] Travers M, Paley S, Shrager J et al. Groups: knowledge spreadsheets for symbolic biocomputing. Database 2013. URL http://database.oxfordjournals.org/content/2013/bat061.abstract.

[46] BioCyc genome and metabolic map posters. http://biocyc.org/posters.shtml.

[47] EcoCyc Cellular Overview. http://ecocyc.org/overviewsWeb/celOv.shtml.

[48] Cellular Overview with Animated Display of Gene Expression Data. http://biocyc.org/ov-expr.shtml.

[49] BioCyc Comparative Genome Browser. http://biocyc.org/ECOLI/NEW-IMAGE?type=LOCUS-POSITION\&object=(EG11024+C1725+TRPA+VC1169)&orgids=(ECOLI+ECOL199310+ECOO157+VCHO).

[50] Edgar R C. MUSCLE: multiple sequence alignment with high accuracy and high throughput. Nucleic Acids Res 32(5):1792–1797, 2004.

[51] Querying Pathway/Genome Databases. http://brg.ai.sri.com/ptools/ptools-resources.html.

[52] RCyc API to Pathway Tools. https://github.com/taltman/RCyc/blob/master/DESCRIPTION.

[53] JavaCyc API to Pathway Tools. http://solgenomics.net/downloads/index.pl.

[54] PerlCyc API to Pathway Tools. http://solgenomics.net/downloads/perlcyc.pl.

[55] PythonCyc API to Pathway Tools. http://bioinformatics.ai.sri.com/ptools/pythoncyc.html.

[56] International S. Pathway Tools User's Guide version 19.0, 2015. Available from SRI International.

[57] Genbank Format. http://www.ncbi.nlm.nih.gov/collab/FT/#7.1.2.

[58] Hucka M, Finney A, Bornstein B J et al. Evolving a lingua franca and associated software infrastructure for computational systems biology: the Systems Biology Markup Language (SBML) project. Syst Biol (Stevenage) 1(1):41–53, 2004.

[59] SBML. http://www.sbml.org/.

[60] BioPAX. http://www.biopax.org/.

[61] et al E D. The BioPAX community standard for pathway data sharing. Nat Biotechnol 28(12):935–942, 2010.

[62] UniProt Sequence annotation (Features). http://www.uniprot.org/help/sequence_annotation.

[63] Ashburner M, Ball C, Blake J et al. Gene Ontology: Tool for the unification of biology. Nature Genetics 25:25–29, 2000.

[64] GO Annotation File Formats. http://geneontology.org/page/go-annotation-file-formats.

[65] Yu N Y, Laird M R, Spencer C et al. PSORTdb–an expanded, auto-updated, user-friendly protein subcellular localization database for Bacteria and Archaea. Nucleic Acids Res 39(Database issue):D241–4, 2011.

[66] Cipriano M J, Novichkov P N, Kazakov A E et al. RegTransBase–a database of regulatory sequences and interactions based on literature: a resource for investigating transcriptional regulation in prokaryotes. BMC Genomics 14:213, 2013.

[67] Barrett T, Wilhite S E, Ledoux P et al. NCBI GEO: archive for functional genomics data sets–update. Nucleic Acids Res 41(Database issue):D991–5, 2013.

[68] Hu J C, Sherlock G, Siegele D A et al. PortEco: a resource for exploring bacterial biology through high-throughput data and analysis tools. Nucleic Acids Res 42(Database issue):D677–84, 2014.

[69] Pathway Tools Data-File Formats. http://bioinformatics.ai.sri.com/ptools/flatfile-format.html.

[70] FASTA Format. http://www.ncbi.nlm.nih.gov/blast/fasta.shtml.

[71] BioCycPlugin for Cytoscape. http://www.cgl.ucsf.edu/cytoscape/bioCycPlugin/.

[72] The BioVelo Query Language. http://biocyc.org/bioveloLanguage.html.

[73] BioVelo Advanced Query Page. http://biocyc.org/query.html.

[74] Pathway Tools Web Services. http://biocyc.org/web-services.shtml.

[75] Latendresse M, Krummenacker M, and Karp P D. Optimal metabolic route search based on atom mappings. Bioinformatics 30:2043–50, 2014.

[76] BioCyc Comparative Analysis. http://biocyc.org/comp-genomics?tables=reaction&tables=pathway&tables=compound&tables=protein&tables=orthologs&tables=tport&tables=tu&savedlist=none&orgid=VCHO&orgid=ECOLI&orglistname=&orglistname=none&orgids=(VCHO+ECOLI+).

[77] Gat E. Point of view: Lisp as an alternative to Java. Intelligence: New Visions of AI in Practice 11(4):21–4, 2000.

[78] Weininger D. SMILES, a chemical language and information system. 1. Introduction to methodology and encoding rules. J Chem Inf Comput Sci 28:31–36, 1988.

[79] Karp P, Lowrance J, Strat T et al. The Grasper-CL Graph Management System. LISP and Symbolic Computation 7:245–282, 1994.

[80] Yan T, Yoo D, Berardini T Z et al. PatMatch: a program for finding patterns in peptide and nucleotide sequences. Nucleic Acids Res 33(Web Server issue):W262–6, 2005.

[81] Bornstein B J, Keating S M, Jouraku A et al. LibSBML: an API library for SBML. Bioinformatics 24(6):880–881, 2008.

[82] Stein S E, Heller S R, and Tchekhovskoi D. An Open Standard for Chemical Structure Representation: The IUPAC Chemical Identifier. In Proc. 2003 International Chemical Information Conference (Nimes), 131–43. 2003.

[83] Page S H. Http://scip.zib.de/.

[84] SKIPPY – Read and Write GIF Files with Common Lisp. http://www.xach.com/lisp/skippy/.

[85] Yahoo User Interface Library. http://developer.yahoo.com/yui/.

[86] Karp P. The design space of frame knowledge representation systems. Technical Report 520, SRI International AI Center, 1992. URL http://www.ai.sri.com/pubs/files/236.pdf.

[87] Paley S and Karp P. Adapting EcoCyc for use on the World Wide Web. Gene 172:GC43–GC50, 1996.

[88] Thorleifsson S G and Thiele I. rBioNet: A COBRA toolbox extension for reconstructing high-quality biochemical networks. Bioinformatics 27(14):2009–10, 2011.

[89] Thiele I and Palsson B O. A protocol for generating a high-quality genome-scale metabolic reconstruction. Nat Protoc 5(1):93–121, 2010.

[90] Harcombe W R, Riehl W J, Dukovski I et al. Metabolic resource allocation in individual microbes determines ecosystem interactions and spatial dynamics. Cell Rep 7(4):1104–1115, 2014.

[91] Latendresse M. Efficiently Gap-Filling Reaction Networks. BMC Bioinformatics 15:225, 2014. URL http://dx.doi.org/10.1186/1471-2105-15-225.

[92] Mueller L, Zhang P, and Rhee S. AraCyc, a biochemical pathway database for Arabidopsis. Plant Physiology 132:453–460, 2003.

[93] AraCyc Database. http://www.arabidopsis.org/biocyc/.

[94] Seo S and Lewin H A. Reconstruction of metabolic pathways for the cattle genome. BMC Systems Biology 3:33, 2009.

[95] BioCyc Database Collection. http://BioCyc.org/.

[96] FlyCyc Database. http://biocyc.org/FLY/organism-summary?object=FLY.

[97] Romero P, Wagg J, Green M et al. Computational Prediction of Human Metabolic Pathways from the Complete Human Genome. Genome Biology 6(1):1–17, 2004. URL http://genomebiology.com/2004/6/1/R2.

[98] HumanCyc Database. http://HumanCyc.org/.

[99] MouseCyc Database. http://mousecyc.jax.org:8000/.

[100] CoffeaCyc Database. http://solcyc.solgenomics.net/organism-summary?object=COFFEA.

[101] Urbanczyk-Wochniak E and Sumner L W. MedicCyc: A biochemical pathway database for Medicago truncatula. Bioinformatics 23(11):1418–1423, 2007.

[102] MediCyc Web Site. http://www.noble.org/mediccyc/.

[103] AraCycGramene Database. http://pathway.iplantcollaborative.org/organism-summary?object=ARA.

[104] BrachyCyc Database. http://pathway.iplantcollaborative.org/organism-summary?object=BRACHY.

[105] CoffeaCycGramene Database. http://pathway.iplantcollaborative.org/organism-summary?object=COFFEA.

[106] MediCycGramene Database. http://pathway.iplantcollaborative.org/organism-summary?object=MEDIC.

[107] RiceCyc Web Site. http://pathway.gramene.org/gramene/ricecyc.shtml.

[108] PoplarCycGramene Database. http://pathway.iplantcollaborative.org/organism-summary?object=POPLAR.

[109] SorghumCyc Database. http://pathway.iplantcollaborative.org/organism-summary?object=SORGHUM.

[110] LycoCycGramene Database. http://pathway.iplantcollaborative.org/organism-summary?object=LYCO.

[111] PotatoCycGramene Database. http://pathway.iplantcollaborative.org/organism-summary?object=POTATO.

[112] MaizeCyc Database. http://pathway.iplantcollaborative.org/organism-summary?object=MAIZE.

[113] SolCyc Web Site. http://sgn.cornell.edu/tools/solcyc/.

[114] CapCyc Database. http://solcyc.solgenomics.net/organism-summary?object=CAP.

[115] BenthmianaCyc Database. http://solcyc.solgenomics.net/organism-summary?object=BENTH.

[116] TobaccoCyc Database. http://solcyc.solgenomics.net/organism-summary?object=TOBACCO.

[117] PetCyc Database. http://solcyc.solgenomics.net/organism-summary?object=PET.

[118] LycoCyc Database. http://solcyc.solgenomics.net/organism-summary?object=LYCO.

[119] PotatoCyc Database. http://solcyc.solgenomics.net/organism-summary?object=POTATO.

[120] PlantCyc Website. http://plantcyc.org/.

[121] BrachypodiumCyc Database. http://pmn.plantcyc.org/organism-summary?object=BRACHYPODIUM.

[122] ChineseCabbageCyc Database. http://pmn.plantcyc.org/organism-summary?object=CHINESECABBAGE.

[123] PapayaCyc Database. http://pmn.plantcyc.org/organism-summary?object=PAPAYA.

[124] ChlamyCyc Database. http://pmn.plantcyc.org/organism-summary?object=CHLAMY.

[125] CassavaCyc Database. http://pmn.plantcyc.org/organism-summary?object=CASSAVA.

[126] BarleyCyc Database. http://pmn.plantcyc.org/organism-summary?object=BARLEY.

[127] SoyCyc Database. http://pmn.plantcyc.org/organism-summary?object=SOY.

[128] SwitchgrassCyc Database. http://pmn.plantcyc.org/organism-summary?object=SWITCHGRASS.

[129] MossCyc Database. http://pmn.plantcyc.org/organism-summary?object=MOSS.

[130] PoplarCyc Database. http://pmn.plantcyc.org/organism-summary?object=POPLAR.

[131] OryzaCyc Database. http://pmn.plantcyc.org/organism-summary?object=ORYZA.

[132] SelaginellaCyc Database. http://pmn.plantcyc.org/organism-summary?object=SELAGINELLA.

[133] SetariaCyc Database. http://pmn.plantcyc.org/organism-summary?object=Setaria.

[134] SorghumBicolor Database. http://pmn.plantcyc.org/organism-summary?object=SORGHUMBICOLOR.

[135] GrapeCyc Database. http://pmn.plantcyc.org/organism-summary?object=GRAPE.

[136] CornCyc Database. http://pmn.plantcyc.org/organism-summary?object=CORN.

[137] Ano2Cyc Database. http://biocyc.org/organism-summary?object=ANO2.

[138] CandidaCyc Database. http://pathway.candidagenome.org/.

[139] DictyCyc Web Site. http://dictybase.org/Dicty_Info/dictycyc_info.html.

[140] LeishCyc Web Site. http://bioinformatics.bio21.unimelb.edu.au/leishcyc.html.

[141] PchrCyc Database. http://biocyc.org/organism-summary?object=PCHR.

[142] YeastCyc Database. http://pathway.yeastgenome.org/.

[143] ThapsCyc Database. http://biocyc.org/organism-summary?object=THAPS.

[144] TrypanoCyc Database. http://vm-trypanocyc.toulouse.inra.fr/TRYPANO/organism-summary.

[145] EcoCyc Database. http://EcoCyc.org/.

[146] BsubCyc Database. http://BsubCyc.org/.

[147] Cbtq1Cyc Database. http://biocyc.org/organism-summary?object=CBTQ1.

[148] PabtqvlcCyc Database. http://biocyc.org/organism-summary?object=PABTQVLC.

[149] EvaCyc Database. http://biocyc.org/organism-summary?object=EVA.

[150] ClossacCyc Database. http://biocyc.org/organism-summary?object=CLOSSAC.

[151] 10403s_rastCyc Database. http://biocyc.org/organism-summary?object=10403S_RAST.

[152] Mob3bCyc Database. http://biocyc.org/organism-summary?object=MOB3B.

[153] Pdif272563Cyc Database. http://biocyc.org/organism-summary?object=Pdif272563.

[154] ScoCyc Web Site. http://scocyc.streptomyces.org.uk:14980/SCO/.

[155] PseudoCyc Database. http://v2.pseudomonas.org:1555/.

[156] Page M H. Https://www.genoscope.cns.fr/agc/microscope/metabolism/microcyc.php.

[157] TBestDB: Taxonomically Broad EST Database. http://tbestdb.bcm.umontreal.ca/searches/welcome.php.

[158] O'Brien E A, Koski L B, Zhang Y et al. TBestDB: a taxonomically broad database of expressed sequence tags (ESTs). Nuc Acids Res 35(1):D445–451, 2007. URL http://www.ncbi.nlm.nih.gov/pmc/articles/PMC1899108/.

[159] BeoCyc Database. http://cricket.ornl.gov/cgi-bin/beocyc_home.cgi.

[160] Pathway Tools PGDB Registry. http://BioCyc.org/registry.html.

[161] Kanehisa M, Goto S, Hattori M et al. From genomics to chemical genomics: New developments in KEGG. Nuc Acids Res 34:D354–D357, 2006.

[162] Okuda S, Yamada T, Hamajima M et al. KEGG Atlas mapping for global analysis of metabolic pathways. Nuc Acids Res 36:W423–26, 2008.

[163] Overbeek R, Olson R, Pusch G D et al. The SEED and the Rapid Annotation of microbial genomes using Subsystems Technology (RAST). Nucleic Acids Res 42(Database issue):D206–D214, 2014.

[164] Salomonis N, Hanspers K, Zambon A C et al. GenMAPP 2: New features and resources for pathway analysis. BMC Bioinformatics 8:217–28, 2007.

[165] van Iersel M P, Kelder T, Pico A R et al. Presenting and exploring biological pathways with Pathvisio. BMC Bioinformatics 9:399, 2008.

[166] Krishnamurthy L, Nadeau J, Ozsoyoglu G et al. Pathways database system: An integrated system for biological pathways. Bioinformatics 19(8):930–937, 2003.

[167] Ozsoyoglu Z M, Ozsoyoglu G, and Nadeau J. Genomic pathways database and biological data management. Anim Genet 37 Suppl 1:41–47, 2006.

[168] Hu Z, Ng D M, Yamada T et al. VisANT 3.0: New modules for pathway visualization, editing, prediction and construction. Nuc Acids Res 35:W625–32, 2007.

[169] Croft D, Mundo A F, Haw R et al. The Reactome pathway knowledgebase. Nucleic Acids Res 42(Database issue):D472–7, 2014.

[170] Shannon P, Markiel A, Ozier O et al. Cytoscape: A software environment for integrated models of biomolecular interaction networks. Genome Research 13(11):2498–2504, 2003.

[171] Klukas C and Schreiber F. Dynamic exploration and editing of KEGG pathway diagrams. Bioinformatics 23(3):344–350, 2007.

[172] Okuda S, Yamada T, Hamajima M et al. KEGG Atlas mapping for global analysis of metabolic pathways. Nuc Acids Res 36(suppl 2):W423–6, 2008. URL http://nar.oxfordjournals.org/content/36/suppl_2/W423.abstract.

[173] Karp P, Krummenacker M, Paley S et al. Integrated pathway/genome databases and their role in drug discovery. Trends in Biotechnology 17(7):275–81, 1999.

[174] Kono N, Arakawa K, Ogawa R et al. Pathway Projector: Web-Based Zoomable Pathway Browser Using KEGG Atlas and Google Maps API. PLoS ONE 4(11):e7710, 2009. URL http://dx.doi.org/10.1371%2Fjournal.pone.0007710.

[175] Stein L D, Mungall C, Shu S et al. The generic genome browser: A building block for a model organism system database. Genome Research 12(10):1599–1610, 2002.

[176] Donlin M J. Using the Generic Genome Browser (GBrowse). Curr Protoc Bioinformatics Chapter 9:Unit 9.9, 2007.

[177] Skinner M E, Uzilov A V, Stein L D et al. JBrowse: a next-generation genome browser. Genome Research 19(9):1630–1638, 2009.

[178] Markowitz V M, Chen I M, Palaniappan K et al. IMG 4 version of the integrated microbial genomes comparative analysis system. Nucleic Acids Res 42(Database issue):D560–D567, 2014.

[179] Entrez Genome. http://www.ncbi.nlm.nih.gov/genome?db=genome.

[180] Rosenbloom K R, Armstrong J, Barber G P et al. The UCSC Genome Browser database: 2015 update. Nucleic Acids Res 43(Database issue):D670–81, 2015.

[181] Cunningham F, Amode M R, Barrell D et al. Ensembl 2015. Nucleic Acids Res 43(Database issue):D662–9, 2015.

[182] Wattam A R, Abraham D, Dalay O et al. PATRIC, the bacterial bioinformatics database and analysis resource. Nucleic Acids Res 42(Database issue):D581–D591, 2014.

[183] Wang H, Su Y, Mackey A J et al. SynView: a GBrowse-compatible approach to visualizing comparative genome data. Bioinformatics 22(18):2308–2309, 2006.

[184] McKay S J, Vergara I A, and Stajich J E. Using the Generic Synteny Browser (GBrowse_syn). Curr Protoc Bioinformatics Chapter 9:Unit 9.12, 2010.

[185] Kazakov A E, Cipriano M J, Novichkov P S et al. RegTransBase—A database of regulatory sequences and interactions in a wide range of prokaryotic genomes. Nuc Acids Res 35(Database issue):D407–D412, 2007.

[186] Matys V, Kel-Margoulis O V, Fricke E et al. TRANSFAC and its module TRANSCompel: Transcriptional gene regulation in eukaryotes. Nuc Acids Res 34(Database issue):D108–D110, 2006.

[187] Baumbach J. CoryneRegNet 4.0 — A reference database for corynebacterial gene regulatory networks. BMC Bioinformatics 8:429, 2007.

[188] Klein J, Leupold S, Munch R et al. ProdoNet: Identification and visualization of prokaryotic gene regulatory and metabolic networks. Nuc Acids Res 36(Web Server issue):W460–4, 2008.

[189] Sierro N, Makita Y, de Hoon M et al. DBTBS: a database of transcriptional regulation in *Bacillus subtilis* containing upstream intergenic conservation information. Nuc Acids Res 36(Database issue):D93–6, 2008.

[190] Gama-Castro S, Jiménez-Jacinto V, Peralta-Gil M et al. RegulonDB (version 6.0): Gene regulation model of *Escherichia coli* K–12 beyond transcription, active (experimental) annotated promoters and Textpresso navigation. Nuc Acids Res 36:D120–4, 2008.

[191] Drysdale R A, Crosby M A, Gelbart W et al. FlyBase: Genes and gene models. Nuc Acids Res 33:D390–D395, 2005.

[192] Bult C J, Eppig J T, Kadin J A et al. The Mouse Genome Database (MGD): Mouse biology and model systems. Nuc Acids Res 36:D724–D728, 2008.

[193] Aurrecoechea C, Brestelli J, Brunk B P et al. PlasmoDB: A functional genomic database for malaria parasites. Nuc Acids Res 37(Database issue):D539–43, 2009.

[194] Smedley D, Haider S, Ballester B et al. BioMart — biological queries made easy. BMC Genomics 10:22–33, 2009.

[195] Birkland A and Yona G. Biozon: A system for unification, management and analysis of heterogeneous biological data. BMC Bioinformatics 7:70–93, 2006.

[196] Dahlquist K, Salomonis N, Vranizan K et al. GenMAPP, a new tool for viewing and analyzing microarray data on biological pathways. Nature Genetics 31:19–20, 2002.

[197] Holford M, Li N, Nadkarni P et al. VitaPad: Visualization tools for the analysis of pathway data. Bioinformatics 15:1596–1602, 2004.

[198] Chung H, Kim M, Park C et al. ArrayXPath: Mapping and visualizing microarray gene-expression data with integrated biological pathway resources using Scalable Vector Graphics. Nuc Acids Res 32:W460–4, 2004.

[199] Henry C S, DeJongh M, Best A A et al. High-throughput generation, optimization and analysis of genome-scale metabolic models. Nat Biotechnol 28(9):977–82, 2010.

[200] Lakshmanan M, Koh G, Chung B K et al. Software applications for flux balance analysis. Brief Bioinform 15(1):108–22, 2014.

[201] Hamilton J J and Reed J L. Software platforms to facilitate reconstructing genome-scale metabolic networks. Environ Microbiol 16(1):49–59, 2014.

[202] Jeong H, Tombor B, Albert R et al. The large-scale organization of metabolic networks. Nature 407:651–4, 2000.

[203] Ma H W, Kumar B, Ditges U et al. An extended transcriptional regulatory network of *Escherichia coli* and analysis of its hierarchical structure and network motifs. Nuc Acids Res 32(22):6643–6649, 2004.